Workshop on the Economics of Information Security

Can We Afford Integrity by Proof-of-Work? Scenarios Inspired by the Bitcoin Currency

Jörg Becker, Dominic Breuker[1], Tobias Heide,
Justus Holler, Hans Peter Rauer, and Rainer Böhme

Department of Information Systems
University of Münster - ERCIS, Germany

{becker,breuker,heide,holler,rauer,boehme}@ercis.de

Abstract

Proof-of-Work (PoW), a well-known principle to ration resource access in client-server relations, is about to experience a renaissance as a mechanism to protect the integrity of a global state in distributed transaction systems under decentralized control. Most prominently, the Bitcoin cryptographic currency protocol leverages PoW to 1) prevent double spending and 2) establish scarcity, two essential properties of any electronic currency. This paper asks the important question whether this approach is generally viable. Citing actual data, it provides a first cut of an answer by estimating the resource requirements, in terms of operating cost and ecological footprint, of a suitably dimensioned PoW infrastructure and comparing them to three attack scenarios. The analysis is inspired by Bitcoin, but generalizes to potential successors, which fix Bitcoin's technical and economic teething troubles discussed in the literature.

1 Introduction

Proof-of-Work (PoW) is a principle to artificially impose transaction costs in the absence of a payment system. The idea is to "charge" the requester of a service with the effort to present a solution to a problem that is much harder to solve than to verify. This way, PoW can help to ration access to services which would otherwise be abused. Originally presented by Dwork & Naor (1992) as a mechanism to combat junk email, it has been proposed as a solution for numerous other situations in which the goal is to prevent some sort of fraudulent use, e.g., when measuring the number of visitors a website has (Franklin & Malkhi, 1997). Whether PoW actually has any practical relevance has been debated in the literature, mainly for the case of spam prevention (see (Laurie & Clayton, 2004) and (Liu & Camp, 2006)).

A particularly innovative use of this principle has been proposed by Nakamoto (2008), who used it as a core component in designing a fully decentralized peer-to-peer electronic currency system called Bitcoin. In this system, users constantly participate in a lottery, and each user's chance of winning is proportional to the computing power he is willing to invest. The task is to modify a document until its hash is of a particular structure. In parallel, users publicly announce transactions of Bitcoins, thereby expressing their intention to transfer a certain amount of this currency to another user. The lottery is designed in such a way that every win, as a side-effect, returns a timestamp of all transactions in an atomic operation. Furthermore, the user who won the lottery is rewarded for his effort by receiving new Bitcoins as well as fees for timestamped transactions. Other users accept this timestamp after validating that the PoW has been delivered. Then, the next round of the lottery starts. The sequence of timestamps forms a history of transactions. In case that competing histories emerge (which is easily possible in a peer-to-peer network), users believe in the history for which the most PoW has been delivered. This quickly resolves the conflict. Altering any transaction timestamped in the past requires redoing all work that has been done afterwards.

[1] Corresponding author

As this becomes more unlikely the longer a transaction has been timestamped, manipulation is prevented and users collectively agree on a single history of transactions, thereby determining Bitcoin ownership.

As an advantage of such a decentralized currency, Nakamoto points out that electronic payment nowadays heavily relies on central authorities processing them. Widely trusted (but not necessarily trustworthy) financial institutions handle electronic payments and ensure the integrity of the system's global state. In return, they charge society for this service. The goal of Bitcoin is to replace trust in financial institutions with trust in PoW, thereby eliminating the need for financial institutions and with it the fees they charge. However, running the Bitcoin system is not for free either. As computational power is required for the PoW, hardware has to be acquired, powered, and maintained.

The Bitcoin system is secure as long as no single party controls more than 50% of the network's computing power. If a party would, it could redo work of the past and ultimately outperform the honest part of the network. While digital signatures on transactions prevent arbitrary manipulations, the attackers would have the power to double-spend their own Bitcoins, thereby generating profits for themselves, or to prevent transactions from being timestamped, thereby undermining the trust in this system. The viability of the system depends on the assumption that nobody will ever gain this power.

These considerations suggest that the overall computational power required to run a system like Bitcoin depends on security considerations. If computational power would be low enough to allow a single party to gain control over 50% of it no one could trust in PoW anymore. Consequently, the cost of running the system depends on security requirements, as computing power drives costs. The aim of this paper is to propose an estimation of what the cost-saving potential of an electronic currency relying on PoW could be. To accomplish this, we present three different attack scenarios. In each of them, an individual or a group acquires control over computing power with the purpose of compromising the integrity of the system. Estimating their computational power delivers an idea of what we would need to defend the system. As a point of reference, we also estimate the transaction costs that a central electronic payment system would produce if used globally, and imagine a distributed network for PoW generation of equal costs. Comparing this network with the attackers allows us to estimate by how much the network could be downscaled securely.

Another interesting aspect is the environmental impact of large-scale PoW application. As electricity costs constitute a large part of the total costs of providing computing power, running a PoW-based currency system would consume a considerable amount of power. In turn, this means it could be responsible for a substantial amount of carbon dioxide (CO_2) emissions and may contribute to global warming. Therefore, we make an interesting detour in estimating the CO_2 footprint of a global PoW-based electronic currency.

The remainder of this paper is structured as follows. Section two discusses details of the Bitcoin system, in particular the role of PoW, and explains how an attack on a currency relying on PoW can be accomplished. In section three, previous research regarding Bitcoin is surveyed. Section four then presents our proposal for estimating potential cost savings as well as the environmental impact of PoW-based currencies. Results from this estimation as well as its limitations are discussed in section five. Finally, section six concludes and provides an outlook on future research.

2 The Bitcoin System

Possessing a certain amount of a currency means possessing the promise that one can collect favors from others in the future whose value is equal to those oneself had to give to others to acquire the amount. Thus, scarcity is a necessary property of anything that is used as a currency. If it could be produced at low cost, devaluation would destroy the trust in this promise. Building upon this property, it is also necessary that the currency can be transferred from one person to another, ideally in any quantity. The transaction mechanism must make sure the currency is correctly transferred in the sense that the overall amount of currency is preserved and that the transaction cannot be reversed later on. We will call this integrity.

For ordinary physical cash, both scarcity and integrity are enforced by the laws of nature, i.e., usage of physical tokens (also, legislation imposes constraints on counterfeit money, theft, …). With electronic currencies however, the amount of currency one possesses is nothing but an account balance. Thus, it requires a financial institution to manage these accounts and enforce scarcity and integrity. Bitcoin (Nakamoto, 2008) has been proposed as a distributed peer-to-peer accounting system accomplishing this without relying on trust in a central authority.

A Bitcoin is represented by a chain of publicly announced transactions every participant of the network is aware of. Each transaction contains a public key signifying the owner of the Bitcoin. It further contains a hash of both the previous transaction and this public key, thereby entangling the previous transaction with the next one. The previous owner signs this hash using his private key. Any user can now validate if the signature of transaction matches the public key of the previous owner. He will only accept the transaction if it does. Thus, knowledge of the private key enables a user to spend a Bitcoin.

As long as users keep their private keys secret, this mechanism prevents potential attackers from spending Bitcoins they do not own. Nothing however stops them from spending those they once received twice. What is required to prevent double spending is a consensus on a temporal ordering of transactions among all users. This way, the current owner of a Bitcoin can always be determined. Attempts of previous owners to spend it again can be detected.

The temporal order is established by what is called a block chain. Blocks collect transactions combined with a hash of the previous block, thereby creating a chain. A block is valid only if it exhibits a special property which proves that a certain amount of work has been put into its creation. In particular, blocks contain a nonce value. It must be chosen by the creator in such a way that, when hashing the block, the hash starts with a certain number of zeros. As for an ideal hash function, this can only be achieved by randomly trying many different nonce values. The probability of finding a block depends on the number of trials. It can be adjusted globally by changing the number of leading zeros the hash must have. Regular adjustments ensure that new blocks are found on average each ten minutes.

Via the block chain users collectively agree on the temporal order of transactions as defined by the order of the blocks containing them. Users always believe in the longest valid block chain they are aware of. As long as the majority of users are honest, no single attacker will be in the position to deliver the PoW necessary to change the temporal order to his advantage. Removing a transaction one did an hour ago (i.e., about 5-6 blocks in the past) from the block containing it would not only require recreating this block but also all subsequent ones, since the hash of the altered block is part of the next, and so forth. Thus, it quickly becomes computationally intractable to alter blocks deep within the chain.

To motivate users to participate in creating blocks, optional transaction fees can be paid by users creating transactions to users ironing these into a new block. Furthermore, a special transaction rewards the creator with a certain amount of newly created Bitcoins (in form of a special initial transaction preceding every valid Bitcoin). This rewarding mechanism—called Bitcoin mining—constantly issues new currency. The amount is reduced at regular intervals and will eventually stop by convention. Once stopped, all Bitcoins are in circulation and the system will solely be in the transaction phase (as opposed to the mining phase in which new Bitcoins are created). Consequently, scarcity of Bitcoins is ensured by eventually stopping Bitcoin supply, and integrity is ensured by digitally signed transactions timestamped in a publicly visible block chain receiving credibility by the work invested into its creation.

As pointed out before, the underlying assumption is that the majority of users (in terms of computing power) are honest, which appears reasonable in a large distributed peer-to-peer network. If however a single user gains control over the majority of the (distributed) computing power, he is able to manipulate the temporal order of transactions to his advantage. The attack is to spend a Bitcoin, which then gets incorporated into the honest block chain. This makes the recipient believe he now possesses the coin. The attacker secretly computes a second chain not containing this transaction and outpaces the honest one. Once he publishes it, users will believe in the attacker's chain and he can spent the coin

again, as it now appears to be still his. Scròfina (2011) discusses several possibilities of double-spending and estimates their payoffs based on the Bitcoin system as it currently is.

It is important to notice that the motivation for attacking the system is not necessarily to profit from the attack. It might as well be to purpose of an attack to simply destroy the system, e.g., as a form of terrorism. In case that a destructive attack is launched, an attacker could for instance prevent any transaction from being timestamped. Effectively, no user could be sure anymore that he actually received a Bitcoin in a transaction, which would quickly destroy trust into the currency and devaluate it. Any of the attack scenarios we analyze in this paper is potentially destructive, i.e., we do not require the attack to be profitable.

3 Related Work

Although the idea of cryptographic electronic currencies came up more than two decades ago (Levy, 1994), it took until today for one of them to be widely discussed in media and research. With Digicash (Chaum, Fiat, & Naor, 1988), a digital currency was presented in 1990. A cryptographic protocol allowed for anonymous payments and copying money was impossible. In contrast to the Bitcoin system, it relied on a central authority. While gaining quite some attention, it never had a significant breakthrough. Several other attempts to establish electronic currencies, e.g., E-Gold (Zetter, 2009)[2], met a similar fate.

In 2008, the blueprint for the Bitcoin system, a decentralized cryptographic currency, was published (Nakamoto, 2008). The system draws on the principles of b-money (Dai, 1998) and enhances the security of previous electronic currencies with the PoW-based timestamping mechanism. Apart from presumably low transaction costs, a heavily advertised advantage of Bitcoin is that no financial institution has the power to exercise control over money creation or transactions.

While Bitcoin is “hyped” by some, there are several downsides to it as well. Most consumers in e-commerce prefer prices in their local currency (Grinberg, 2011). With the heavily fluctuating exchange rates currently observed, it is hard to price goods (Grigg, 2011). Concerns regarding massive deflation have been expressed as well (Krugman, 2011). Yet this deflation also creates incentives for early adopters to promote the currency in order to profit from changes in exchange rates later on. Hence, some critics describe Bitcoin as a Ponzi scheme which over-rewards early adopters (Barok, 2011; Grigg, 2011). Bitcoin is also a topic for legal scholars, who ask the question how law should deal with such a currency (see e.g., Grinberg (2011)). Being a cryptographic currency using a public/private key mechanism, anonymity is a feature often promised, but contested in the literature (Grinberg, 2011; Reid & Harrigan, 2011). In fact, unlike Digicash, Bitcoin has never been designed to be anonymous.

As for security, there are numerous threats to Bitcoin users other than double-spending. For instance, losing one’s private keys due to computer malware means losing one’s Bitcoins (Chirgwin, 2011; Doherty, 2011; Grinberg, 2011). Also, denial-of-service attacks on the underlying peer-to-peer communication infrastructure or the Bitcoin exchanges pose a threat (Grinberg, 2011; Nakamoto, 2008).

Not only is the current technical implementation of the system subject to critique, but also the protocol itself. Babaioff, Dobzinski, Oren, & Zohar (2011) claim that nodes in the Bitcoin network have an incentive not to propagate transactions. They also propose a modification to solve this problem. Several other problems have been identified (Barber, Boyen, Shi, & Uzun, 2012).

However, any threats or problems discussed above are specific to the current implementation of the Bitcoin protocol, not to the idea of using a PoW-based timestamping mechanism to ensure integrity of an electronic currency. To the best of our knowledge, no scholar has yet analyzed the economic or ecologic consequences of a large-scale application of this mechanism. Therefore, we present a hypothetical scenario of a PoW-

[2] Ultimately the E-Gold company was pled guilty in money laundering and in operating an unlicensed money-transmitting business (US District Court for the District of Columbia, 2006) and was shut down immediately.

based currency system in the next chapter. The only threat we consider is that it is designed to defend against: a single party gaining control over 50% of the total computation power.

4 Costs of a PoW-based Currency

4.1 General Scenario

To compare the two systems, a PoW-based, Bitcoin-like currency and a centralized solution, we now consider the following scenario. First, we conduct the analysis as if the two systems would be in use today. Thus, the scenario is based on technology currently available. We do not worry about transition problems or speculate about future states of the world, but compare steady states which are, albeit imaginary, inspired by the world as it is. Consequently, we assume that the temporary mining phase is over (or never existed, to counter the Ponzi critique), i.e., there is no fixed reward for creating a new block of the chain. The only reward stems from fees paid by those who carry out transactions.

Second, we assume that payments are handled using a single currency for all transactions worldwide. As a PoW-based currency obviously benefits from economies of scale, we chose to consider the largest conceivable network.

The most important assumption in our analysis is the system to which we compare the PoW network. The main goal of our analysis is to explore the effect of the replacement of a centralized institution by a PoW network. For this reason, it is necessary to compare the PoW network to a centralized system that is, apart from being centralized, as similar as possible. We have chosen a card payment system as it processes payments electronically, just as the PoW network does. Other costs, for instance for printing and distributing cash money, are certainly incurred in today's financial system, yet they would distort the analysis as they are not the result of using a centralized system but of demand for non-electronic payments. We also do not assume that financial intermediation vanishes completely as Bitcoin is no replacement for capital markets. It is just an instrument for payments.

In summary, we compare a rolled-out PoW network to a card payment system, assume that both systems process all ordinary payment transactions of the entire world, and we do not rely on forecasts but on financial estimates and technology of the present.

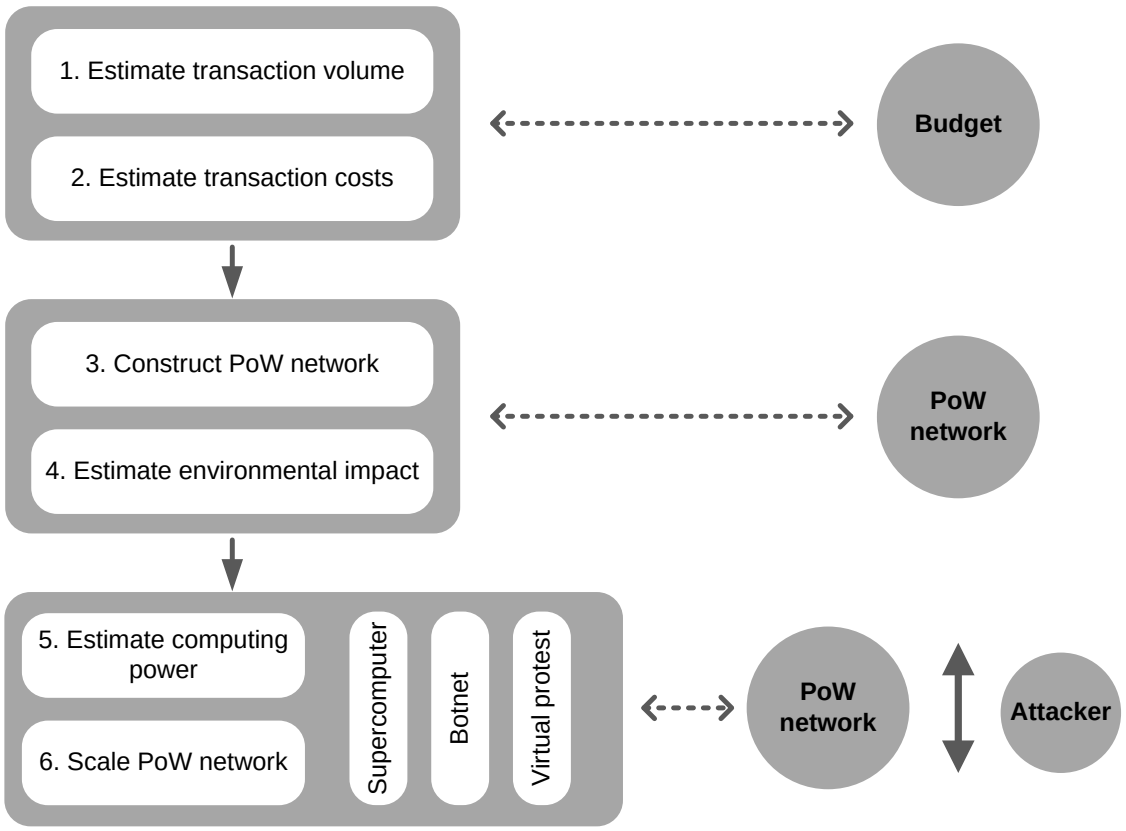

Figure 1: Outline of the analysis

4.2 Analysis Procedure

The idea of our analysis is outlined in figure 1 and structured as follows. First, we will estimate the volume of all transactions taking place anywhere in the world. Second, we estimate the transaction costs incurred in a system using a centralized currency as a fraction of the volume. Combining these two estimates, we end up with the worldwide total transaction costs of a centralized currency. It will serve as a reference budget for the remainder of the calculation. Third, we will design a PoW network using this reference budget, thereby dimensioning a PoW network that costs exactly as much as the centralized system. We will describe this network in terms of the computing power it could achieve. To keep it running, electricity is required. The generation of this electricity has an environmental impact, which is what we will estimate in a fourth step. In step five the computing power of potential attackers is estimated. We consider three different attack scenarios. In the sixth step, given the computing power of both the PoW network and the attacker, we can see if and to which extent the size of the PoW-network could be decreased such that it can still fight off the attack. Finally, this allows estimating the potential for decreasing transaction costs compared to a centralized solution. Additionally, the environmental impact can be estimated.[3]

4.3 Estimate of Transaction Volume

The Bank of International Settlements (BIS), located in Switzerland, publishes annual statistics on payment, clearing, and settlement systems. The reports provide data for 23 countries, among them the United States, Germany, China, Japan, France, the United Kingdom, Russia, Italy, Canada. We will use the transactions taking place in these countries as a proxy for those of the entire world. This appears justified as the economically strongest countries of the world are included. In the latest report (Bank for International Settlements, 2011), the BIS reports the total yearly transaction volume for all 23 countries, categorized by method of payment. Figures for 2010 can be found in table 1.

Table 1: Total value of transactions by payment instrument in US dollars, 2010. Source: (Bank for International Settlements, 2011), comparative table 9

Payment instrument	Transaction volume [USD]
Credit transfers	3.74E+14
Direct debit	4.37E+13
Check	9.36E+13
E-Money	1.80E+10
Card payments (debit/credit)	9.45E+12

Recall that the centralized system to which we compare the PoW network is an electronic card payment system, whose transaction costs will be estimated in step 2 as a certain fraction of the volume. Such systems are used to process small, cash-like transactions. Therefore, large credit transfers between corporations, professional investors, or nations should not be included in the estimate, for applying transaction fees of an electronic card payment system to them would be unreasonable. Consequently, we exclude the item *credit transfers* of table 1 from our estimate. *Direct debit* on the other hand will usually represent cash-like transactions and is therefore included.

The next important item in the list is *check*, which is also a famous payment instrument. Looking at the data, one can see that more than 75% of the check transaction volume stems from China and the United States ((Bank for International Settlements, 2011), comparative table 9). To identify if these transactions are cash-like we have a look at the average volume per transaction ((Bank for International Settlements, 2011), comparative table 9c). For China, it amounts to more than *52000 USD* per transaction, which indicates that the volume stems in large parts from very huge transactions. Thus, we exclude Chinas volume from our estimate. For the United States however, the average value amounts to only

[3] All assumptions of the analysis are debatable. Therefore we provide the spreadsheet of our estimation online and invite readers to come up with their refined scenarios. It can be accessed via: http://dl.dropbox.com/u/1168860/DoesProofOfWorkPayOff.xlsx

3000 USD. This indicates that, while there will surely be large transactions, there will also be a number of small, cash-like ones. Therefore, we include 50% of the check transaction volume of the United States.

All other check transactions are included. This also applies to all *E-Money* transactions as well as *card payments*, delivering a transaction volume of *9.03E+13 USD* in 2010. Transactions for which actual cash is being used are not included yet. For these, we do not know any reliable data source. However, statistics from the BIS report the total volume of ATM cash withdrawals for 2010, which amounts to *4.09E+12 USD* ((Bank for International Settlements, 2011), comparative table 13). Assuming that most of this cash is spent only once and then deposited into a bank account before being withdrawn again, we use this as an estimate for cash transactions. Adding it, we end up with a *total transaction volume of 9.44E+13 USD*.

4.4 Estimation of the Total Transaction Costs of Centralized Payment

Given an estimate of the transaction volume of all cash-like transactions taking place anywhere on the world, we now proceed with estimating the cost of these transactions charged by the financial system to the real economy. As a reference, we will use an electronic payment system, as for such systems market prices are known. In general, there are two different types of them, the first of which is a credit card system. Typically, credit card fees paid by merchants in different countries can vary around 1-3 percent of the transaction volume ((Weiner & Wrigth, 2005), Figure 3). However, pricing of credit cards is subject to considerable suspicion. Retailers have filed numerous lawsuits in the past as they believe institutions issuing these cards exaggerate their costs (Bradford & Hayashi, 2008). Thus, a credit card system is not the ideal reference for comparison with a PoW network. In addition, credit cards are not only used for ordinary payments, but also offer a credit function, i.e., the card holder buys products on credit and pays for them later. This is an additional service not offered by a PoW network like Bitcoin. Naturally, these services charge a credit risk premium. This further strengthens the belief that costs of a credit card system are inadequately high for the purpose of this analysis.

The second type of electronic payment system is a debit card system. In contrast to credit cards, payments made with a debit card are transferred immediately (within days) from the user's bank account to the recipient. It does not provide a credit function. Also, the transaction costs are considerably lower as compared to credit cards. Therefore, a debit card system is used for comparison to the PoW network, as it comes closest to a frictionless centralized payment system free of bells and whistles that have no counterpart in the PoW network. The debit card system used in Germany, called *electronic cash*, charges merchants 0.3% of the transaction volume as a fee (EURO Kartensysteme GmbH, 2008). Applying this to the transaction volume estimated in step 1, we end up *with total transaction costs of 2.83E+11 USD* for an imaginary world in which all transactions are processed with a centralized system comparable to debit cards.

4.5 Size of a PoW Network

In the following step, we use the transaction cost of a centralized system as a budget to size a PoW network. The first and most important question is what the components of that network should be. Today's Bitcoin network is run to a large extent by individuals. They use their PCs and other equipment to create new blocks. However, once block creation becomes a commercial activity, it is unlikely that any kind of equipment owned by individuals could compete with specialized hardware as it is found in data centers. Thus, creating blocks would quickly become unprofitable for them, leaving the market to professional players.

On the other extreme, economies of scale could result in only a very few data centers, owned by a small number of organizations that cover the entire market of PoW delivery. Such a setting would clearly not be the decentralized currency envisioned by the Bitcoin pioneers. However, given that the fate of the global financial system rests on the security of this system, the risk that somebody gained control over these data centers could under no circumstances be acceptable. Moreover, as the organizations running the data centers might collude to exploit the system, one would be forced to trust them. For these reasons, we assume for our scenario a scale between these two extremes. More precisely,

we assume a large number of independent data centers distributed all over the world, each small enough to implement a governance regime with effective checks and balances.

Having defined the main building block of the PoW network, we now analyze the typical cost structure of a data center. According to Belady (2007), there are three main components.

- *Acquisition cost*: The cost of acquiring hardware. Typically amounts to about 25% of the total cost.
- *Energy cost*: The cost of electricity required to run the data center. Typically amounts to about 30% of the total cost.
- *Infrastructure cost*: The cost of providing the environment in which the data center is run (e.g., the rent for the building). Typically amounts to about 45% of the total cost.

Of particular importance for this analysis are the costs for electricity, as these will be the main driver for pollution. We will therefore pursue the following course in the construction of the PoW network. An estimate of computing power is achieved with respect to consumed electricity, under the assumption that particularly energy efficient hardware is used. At the same time, we disregard the fact that such hardware might have higher acquisition cost. This means the data centers of the network are energy efficient but still have the above mentioned cost structure.

It is worth noticing that we do not take into account the costs of the communication infrastructure required to operate the system. For instance, Kaminsky (2011) expresses doubt that the decentralized architecture of the current Bitcoin network could scale up to the size of our PoW network. Other authors already investigate the aspect of scalability and provide suggestions how a large-scale system could be designed (Barber et al., 2012). For our calculation, we assume that smart protocols could overcome all problems. However, as we do not know how such a solution might look like, we do not estimate any communication costs for now and focus solemnly on effort for delivering PoW.

With a total budget of *2.83E+11 USD*, of which 30% are being spent for electricity, we obtain an electricity budget of *8.49E+10 USD.* The next task is to determine the amount of electricity that can be produced. Again, we use current market prices. Typical prices for different countries can be found in table 2, a broad overview in Wikipedia (2012a).

Table 2: Exemplarily electricity prices in various countries of the world

Country	Electricity price [USD/kWh]	Source
Germany	0.36	(Europe's Energy Portal, 2010)
China	0.16	(Yang, 2009)
USA	0.11	(EIA, 2012)
Russia	0.10	(Mosenergosbyt, 2012)

Several countries do have very low electricity prices but are rather unimportant with respect to global energy production. Considering only countries with a major output of electricity, Russia has the lowest (*0.10 USD/kWh*). Consequently, this value is used as a lower bound for the cost of electricity. Given the electricity budget of *8.49E+10 USD*, a total of *8.49E+11 kWh* is available for computation. This equals *3.06E+18 Ws* (watt-seconds) after a unit conversion. Note that this is the amount of electricity available per year, since we have calculated it from annual transaction costs.

We will now estimate the computing power that could be sustained over the year. An important aspect is how computing power should be measured for this purpose. Typically, high performance computing power is measured in floating point operations per second (*FLOPS*). For instance, the Top 500 list, ranking the 500 most powerful supercomputers in the world, exclusively relies on this measure (TOP500.Org, 2012). Delivering PoW in the Bitcoin system however heavily relies on computing hashes, for which integer operations are required[4]. Technically, the *FLOPS* measure should be replaced with a more appropriate metric such as megahash per second. Unfortunately, due to the widespread adoption of

[4] This is specific to the Bitcoin system as it currently is. It is also conceivable to design PoW functions that are better aligned with the optimization criteria of microprocessor architectures.

FLOPS, data on other metrics is not directly available. Even the Bitcoin community itself reports estimates of the network's computing power in terms of FLOPS (Bitcoinwatch.com, 2012). Therefore, we adopt it for this analysis and discuss this as one limitation in section five.

Similar to the Top500 list, the Green500 list reports the 500 most energy efficient supercomputers in the world (Green500.org, 2012) as measured in Mega-FLOPS per watt (MFLOPS/W). In the most recent ranking (November 2011), the most efficient one is an experimental computer from the *IBM Blue Gene/Q* project at *IBM – Rochester*, achieving 2026.48 *MFLOPS/W*. The performance of an "average" green supercomputer however is not even close to this value. In figure 2, a boxplot of the performance distribution in *MFLOPS/W* is presented, with the upper (lower) whisker being the 97.5% (2.5%) quantile.

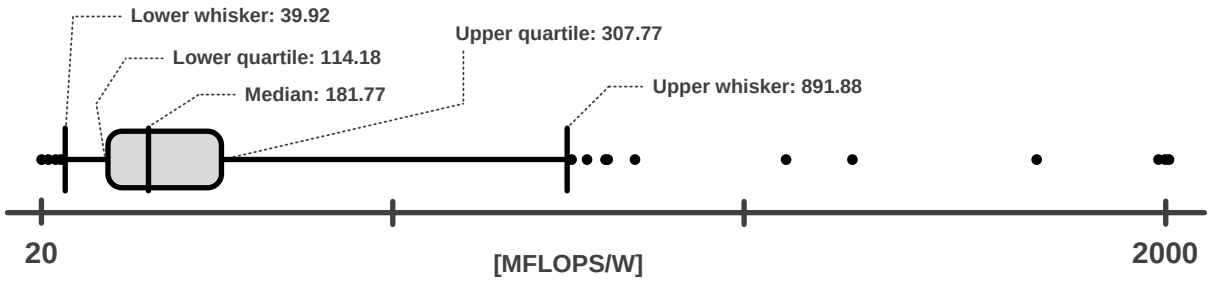

Figure 2: Boxplot of the Green500 November 2011 list

This distribution indicates that most of the computers achieve a performance between 100 and 300 *MFLOPS/W*. For the "average" energy-efficient computer that we use to construct our hypothetical PoW network, we believe that a value of *181.77 MFLOPS/W* is appropriate. It equals the median of the Green500 supercomputers. Note that *MFLOPS/W* means million floating point operations per second per watt, i.e., million floating point operations per watt-second. To keep the notation clear, we will express the efficiency from now on in terms of (floating point) operations per watt-second (*Ops/Ws*).

Multiplying our estimate for energy-efficiency (*1.82E+8 Ops/Ws*) with the available electricity delivers a total *of 5.56E+26 Ops* that can be achieved per year. Divided by the 3*.15E+7* seconds of a year, the *computing power of the PoW network is 1.76E+19 Ops/s.*

4.6 Estimate of Environmental Impact

Given the annual power consumption of *3.06E+18 Ws*, it is now straightforward to estimate the environmental impact of the PoW network. We will measure environmental impact in terms of CO_2 emissions. To start with this, information on how energy is being produced is required. Such data can be obtained from the International Energy Agency, which provides statistics on the worldwide production of electricity in 2009 with respect to the energy carrier (IEA, 2012). The fraction of the total energy production each energy carrier is responsible for can be seen in the upper part of figure 3[5].

Knowing the relative importance of each energy carrier, we can combine this information with corresponding average CO_2 emissions. Figures for that are provided by Lübbert (2007) and can be seen in the lower part of figure 3[6]. Computing an average of the CO_2 emissions, weighted by the relative importance of each carrier, delivers an average CO_2 emission of about *718 gram/kWh* or, after a unit conversion, of *1.99E-7 kg/Ws*.

Multiplying this estimate with the power consumption of *3.06E+18 Ws*, the PoW network is responsible for a total of *6.10E+11 kg* of CO_2 per year. Compared to the total man-made CO_2 emissions due to fuel combustion in 2009, which have been *2.90E+13 kg* (IEA, 2011), the PoW network at this scale would increase CO_2 emissions by more than

[5] The list provided by the IEA (2012) includes an item „hydro" which consists to a large extent out of energy produced by pumped storage plants. As electricity stored in these plants has been generated by other means, we exclude this item. We further exclude items "waste" and "other sources" as their impact on the result is negligible.

[6] For some carriers, Lübbert (2007) provides minimum and maximum estimates for CO_2 emissions. We use their averages in figure 3.

2.1%. This is about the share of global commercial air traffic. We assume the CO_2 emissions of the centralized system to be negligible compared to global emissions, making the effect of the PoW network a net increase.

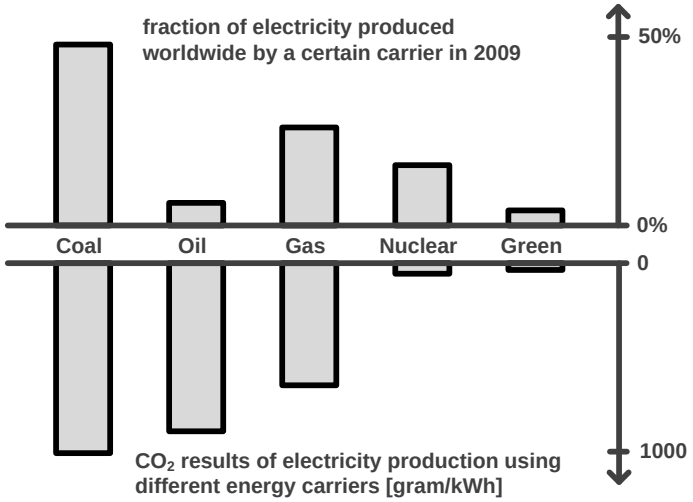

Figure 3: Energy carriers: relative importance in 2009 vs. CO_2 emissions

4.7 Attack by Supercomputer

Whether the size of the constructed PoW network is adequate, too large, or too small can only be judged with respect to the computing power of potential attackers. Once known, the PoW network can be rescaled to the smallest size (with some headroom) such that it is still safe.

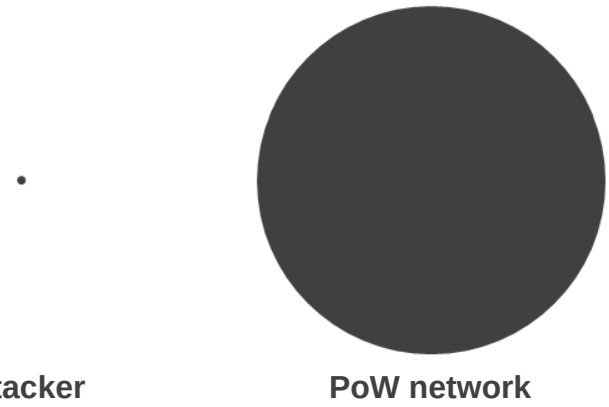

Figure 4: Computing power of supercomputer attacker and PoW network. The area of the dots corresponds to the operations per second.

As a first naïve attack, we compare the network to a large supercomputer. In an extreme scenario, today's largest supercomputer on earth, namely the K computer at Riken in Japan, may attack the system. According to the Top500 list of November 2011, it has a computing power of *1.05E+16 Ops/s* (TOP500.Org, 2012). Compared to the *1.76E+19 Ops/s* of the Pow network, the attacker would only control about 0.06% of the total computing power. Consequently, there is much potential to reduce the size of the network. More precisely, if it would be reduced to 0.12% of its original size, the attacker would be on par with the network. Figure 4 illustrates the computing powers of PoW network and attacker.

4.8 Attack by Botnet

As a next idea, consider an attack in which a large botnet is leveraged to compete with the PoW network. There are mainly two variables that need to be estimated: the size of the botnet and the computing power of an average bot.

Measuring the size of botnets is a topic of active research. In Zhu et al. (2008), several different methods are being reviewed. Rajab, Zarfoss, Monrose, & Terzis (2007) point out the difficulties in generating reliable estimates. For instance, the botnet *Storm*, which has been infiltrated by a security analyst of University College San Diego in 2007, has been estimated to consist of up to 50 million bots (McMillan, 2007), yet the analysis of the security analyst revealed about 200,000 bots being online at a given time and a total of 1.5 million bots per day, thus indicating a considerably smaller size (Enright, 2007). The article *Botnet* on Wikipedia contains a list of the largest botnets currently known, together with estimates of their size (Wikipedia, 2012b). As we want to create a worst-case attack scenario, we assume the size of the attacking botnet equals the largest currently known botnet (called *BredoLab*), regardless of the fact that it might be an overestimation. In numbers, we assume the botnet controls 30 million bots.

Regarding the computing power of an individual bot, we draw an analogy. The Berkeley Open Infrastructure for Network Computing (BOINC) is a software platform designed to enable distributed scientific computing. It is used by numerous projects, among them the famous SETI@home initiative, and publishes statistics on the number of active users as well as total computing power (Boincstats.com, 2012). At the time of writing, a total of 456876 machines generate about *5.64E+15 Ops/s*, which is an average of *1.23E+10 Ops/s* per machine.

Assuming the same average computation power for any of the 30 million bots, the botnet can achieve a total of *3.70E+17 Ops/s*. Compared to the PoW network, the botnet controls approximately 2.06% of the total computing power. Figure 5 illustrates this ratio. The PoW network had to be shrunken to 4% of its original size to allow the botnet to reach 50% of the total power.

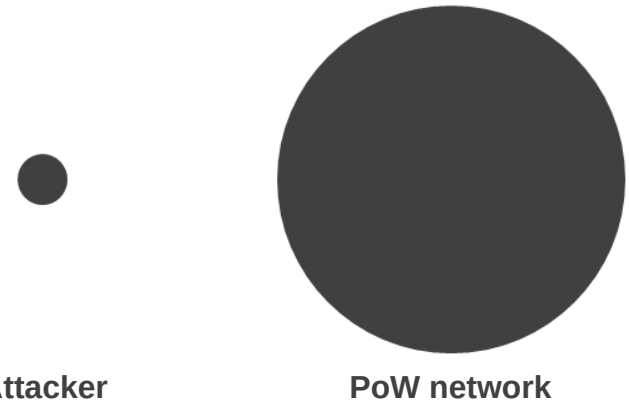

Figure 5: Computing power of botnet attacker and PoW network. The area of the dots corresponds to the operations per second.

4.9 Attack by Virtual Protest ("Occupy Bitcoin")

In the original Bitcoin article, Nakamoto (2008) uses the expression "one-CPU-one-vote" to describe the philosophy of PoW (p. 3). In a world in which all financial transactions – and with them all the world's economies – depend on a PoW network, we want to raise the question: What if a large number of CPU-owners vote against the system?

In particular, consider an attack scenario in which Internet activists acquire a large number of participants via social networks for a virtual protest. In the recent past, the Occupy Wall Street movement found millions of followers. Still, numerous protests are regularly being held all over the world (Wikipedia, 2012c). At October 15 in 2011, global protests, partly inspired by this movement, were held in more than 950 cities with an estimated number of participants between one and two million people (for the 112 locations for which Wikipedia (2012d) provides data). In the protests against the Iraq war in 2003, the total number of participants is estimated to be around 36 million (between January 3rd and April 12th, (Wikipedia, 2012e)), with estimates for the peak value ranging between six and 30 million (on February 15, 2003, (Wikipedia, 2012f)).

Such figures demonstrate that a popular cause can easily unite several million people all over the world and make them cooperate to demonstrate and emphasize their collective opinion. We believe the existence of a PoW network would provide an opportunity for a new form of online virtual protest. If a group of activists provides easy-to-use software and enough participants willing to use it are being found, protesters could try to collaboratively launch a destructive attack against the network. Effectively, this could undermine the trust to an extent such that no transactions are accepted anymore. Unlike ordinary protests, the economic consequences of a successful virtual protest could be devastating. This may be a reason for them being particularly attractive.

Apart from this, virtual protests also lower the effort a protester has to invest for participation. While ordinary protests require traveling to particular location and spending a considerable amount of time there, virtual protests require only downloading and running software. With social networks such as Facebook, the infrastructure to coordinate protesters and distribute software is already in place. For these reasons, we believe that virtual protests could be much bigger than today's ordinary protests ever have been.

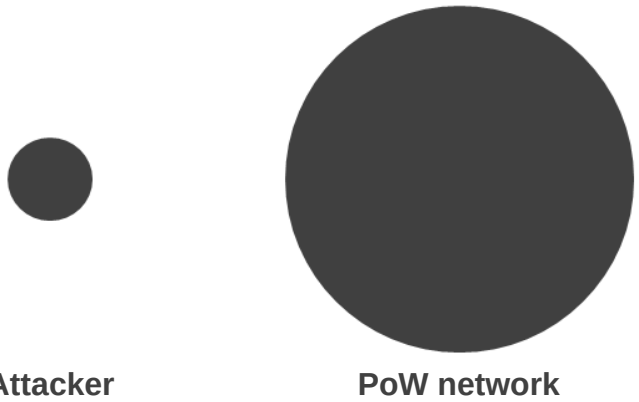

Figure 6: Computing power of protesters and PoW network. The area of the dots corresponds to the operations per second.

To construct a simple scenario, we assume that of the 845 million monthly active Facebook users (Facebook, 2012), 10% are motivated to participate in a virtual protest, e.g., against a war. They all agree to run a suitable piece of software at a particular day. Let the software be designed such that it prevents any transaction from being validated. Further assuming that each participant is the owner of a computer running at the same speed as the bots of our botnet *(1.23E+10 Ops/s),* the attackers' computational power would be *1.04E+18 Ops/s*. This is approximately 5.59% of the total computing power when compared to the PoW network, which is visualized in figure 6. In this scenario, attackers and the network would be on par if the network would be reduced to 12% of its original size.

5 Discussion and Limitations

5.1 Results

Citing actual data and abstracting from frictions specific to the Bitcoin protocol, our analysis sheds light on the true cost of running a decentralized cryptographic currency secured by PoW. Our approach in this paper is to estimate the size of the largest PoW network that could be constructed subject to a cost constraint given by a comparable centralized system, such as one that can process all electronic payments of the developed world today. In addition, we sketch three worst-case attack scenarios, each bounding the potential computing power of attackers from above. This is reasonable because an economy relying on the PoW-based currency as a primary means to process payments could not afford being vulnerable to attack. As a side-effect, our analysis allows us to estimate the ecological footprint of such a PoW-based system.

The results of this estimation exercise suggest that the cost saving potential might be smaller than claimed by Nakamoto (2008) and hoped by the Bitcoin proponents. Even with the large PoW network we considered in our estimation, cutting cost by only one order of (decimal) magnitude would already allow the attackers of the third scenario to overpower the network. Given that one would probably require a decent safety margin over what an attacker might achieve, costs may be cut at best to a fifth of the original amount; notwithstanding that a safety margin of factor two is by no means comparable to typical safety margins in cryptographic security, where the attacker's effort to break a system is calibrated to be at least *1.2E+24* times the defender's effort to use it. Furthermore, a range of political and social obstacles to implementing a PoW-based currency globally further limits the cost saving potential. In particular, a realistic adoption scenario for a PoW network would unlikely be a "big bang" (think of a digital Bretton Woods). Thus, the transition is either very expensive, because the large-scale PoW network must run in parallel to conventional systems, or very risky, as smaller-scale networks remain vulnerable.

In addition, the cost savings would be bought dearly through a substantial increase in global CO_2 emissions. Even if the 2.1% increase of the baseline network would be reduced to a fifth, it would still be an increase of about 0.4% only for direct power consumption. Not yet considered is the environmental impact of producing the required hardware as well as the impact of the communication infrastructure (which is assumed at zero cost and zero emissions in this analysis). In times of growing interest to reduce global emissions, it is questionable if the merits of a PoW-based system justify its environmental cost.

5.2 Robustness

We see our calculation as a first step to understand the longer-term implications of PoW deployment. Of course our results depend on the validity of a number of assumptions. We are the first to admit that some of them are debatable. However, great care has been taken to ensure that our approximations are conservative. To back this up, we list all relevant assumptions in Table 3, state whether they represent upper or lower bounds, and argue why we chose so.

Table 3: Discussion of assumptions

Calculation step	Assumption	Comment
Transaction volume	Upper bound	The largest possible size for the system is that all transactions of the world are processed.
Transaction cost	Upper bound	Using current market prices of centralized electronic payment systems bounds the costs from above as there is no incentive to pay more.
Cost of electricity	Lower bound	Electricity is assumed to be produced at the lowest market price among all countries with major electricity output.
Available electricity	Upper bound	Producing electricity at a lower-bounded price with an upper-bounded budget delivers an upper bound for the available amount.
Cost of computing power	Lower bound	We calculate the cost of computing power in terms of power consumption and assume particularly energy-efficient hardware.
Available computing power	Upper bound	Producing computing power at a lower-bounded price with an upper-bounded budget delivers an upper bound for the available amount.
Environmental impact	Realistic	With respect to the network's power consumption, we believe that CO_2 emissions are fairly realistic given the data that has been used.

Attackers computing power	Upper bound	We consider worst-case attack scenarios with attackers that do not necessarily act economically rational.

5.3 Limitations

Despite the care we have taken in the estimation, some limitations remain. First of all, the measure we use for computing power is FLOPS, while Bitcoin's PoW function is based on hashes. Integer operations or hashes per second would be more appropriate measures. Given that attack and defense are only one magnitude apart, even a small scaling factor could change our conclusions quite a bit. However, other than Bitcoin's hash-based PoW function exist. And once PoW is going to be rolled out in a large infrastructure, demand for hardware optimized to calculate hashes will stimulate R&D and reduce its cost.

To curb the amount of speculation in the analysis, our scenarios are built as counterfactuals assuming that the PoW network is in use today. However, it actually may (or may not) be in use at some time in the future, when technological innovation might have changed the parameters of the scenario. Specifically in the last two of our attack scenarios, a specialized PoW network is compared to attackers who compose their networks largely out of desktop computing hardware. With growing interest in green high performance computing, more energy-efficient specialized hardware might be developed. This would shift the relation between specialized and desktop hardware to the advantage of the PoW network. Moreover, current trends in end-user computing indicate that future consumer devices will be tablets and smartphones instead of desktop computers. Also, functionality is more and more provided through web-based services. Thus, computational power will not be that important for these devices in the future, making it harder to assemble a competing network from them.

Lastly, our analysis ignores the cost of communication because this is more associated with achieving availability rather than integrity. Depending on the design of the underlying peer-to-peer communication network, this cost can be substantial or even prohibitive. Whether it is an advantage or disadvantage for the PoW network is hard to tell, because the two most threatening attack scenarios need extensive communications as well. Therefore we deem it tenable to ignore the costs of communication on both sides for a first and crude analysis.

A string of other limitations stands to reason. We have commented on most of them in the description of the analysis.

6 Conclusions and Future Research

The main goal of this research is to scrutinize the claim of Bitcoin proponents that a decentralized PoW-based currency charges society fewer transaction costs that a centralized electronic payment systems. This is connected to the more general question if society can afford using PoW to enforce the integrity of the global state in a distributed system. The answer to this question has implications on the design and governance of future information infrastructures.

We calculated a conservative cost estimate, accounting for both economic and ecologic costs. If we take our results at face value, they indicate that a PoW network as costly as today's electronic payment systems could easily withstand attacks by a single supercomputer, and most likely defend against a very successful botnet or a social disobedience attack. However, its ecological footprint would be about the share of global commercial air traffic.

The system would still be secure if scaled down somewhat, but it is striking to see that attack and defense are only one (decimal) magnitude apart, although both estimates draw on completely independent inputs and involve plenty of factors. We conclude that although our analysis does not discard PoW networks right away, the question we asked is valid and, given the error margins, it is by no means certain that PoW networks are worth the price they cost.

But this is not the end of the game. Future technological innovations could change the cost-benefit ratio of a PoW network completely, also in favor of PoW. For example, it is

conceivable to reuse byproducts of PoW functions. More specifically, it could be of interest to develop PoW mechanisms that compute something useful. Right now, the activity of finding a nonce value such that the resulting hash satisfies a particular structure is in itself a waste of resources. It is an open research question to formulate relevant problems (e.g., complicated scientific computations, genome sequencing, protein folding) in a form such that a distributed network could solve them and the solution would be easily verifiable. Then anyone who wants a problem solved could formulate it as a PoW function and post a reward on it, thereby financing the PoW network. This could significantly reduce or even nullify the cost of the network and thus contest our result.

Another idea is to reuse the energy instead of (or in addition to) the computation result. Every computation converts electricity into heat pretty efficiently. Therefore, PoW could very well be used to help heat buildings. As decentralization is a key security principle behind PoW networks, a new kind of local compute-heat cogeneration appears much more feasible than transporting the waste heat of large datacenters to the places where it is needed.

Yet another aspect is that the timestamping service provided by the PoW network could serve purposes other than ensuring integrity of a currency system. For instance, Clark & Essex (2011) suggest using it to timestamp commitments in a cryptographic commitment scheme. If enough application scenarios are being found, it might be justified to finance a very large POW-based timestamping service as an infrastructure for all these services. Further research in the above-mentioned areas might eventually allow for a PoW network of an even larger scale than that considered in this paper.

To conclude, we provided a first and very rough proposal to investigate the economic potential of PoW applied to ensure integrity of electronic currency systems. While our results indicate potential for a moderate reduction of economic transaction costs, the ecologic impact is substantial and would surely arouse public resistance if a PoW network were to be established. However, numerous possibilities have been discussed that could turn this result around and make a decentralized PoW timestamping service a valuable infrastructure for the future IT landscape. It is our hope that this paper stimulates both discussions about and further research into these aspects.

References

Babaioff, M., Dobzinski, S., Oren, S., & Zohar, A. (2011). On Bitcoin and Red Balloons. Retrieved from http://arxiv.org/abs/1111.2626

Bank for International Settlements. (2011). Statistics on payment, clearing and settlement systems in the CPSS countries - Figures for 2010. Retrieved from http://www.bis.org/publ/cpss99.htm

Barber, S., Boyen, X., Shi, E., & Uzun, E. (2012). Bitter to Better — How to Make Bitcoin a Better Currency. *16th International Conference on Financial Cryptography and Data Security*. Bonaire, Caribbean Netherlands.

Barok, D. (2011). Bitcoin: censorship-resistant currency and domain system for the people. *Forum American Bar Association*. Retrieved from http://pzwart3.wdka.hro.nl/mediawiki/images/archive/6/64/20110719200925!Barok.bitcoin.pdf

Belady, C. L. (2007). In the data center, power and cooling costs more than the it equipment it support. *Electronics Cooling*, *13*(1), 24-27.

Bitcoinwatch.com. (2012). Bitcoinwatch. Retrieved from http://bitcoinwatch.com/

Boincstats.com. (2012). BOINC Combined Project Statistics. Retrieved February 16, 2012, from http://boincstats.com/stats/project_graph.php?pr=bo

Bradford, T., & Hayashi, F. (2008). Developments in Interchange Fees in the United States and Abroad. *Payments System Research Briefing*. Retrieved from http://www.kc.frb.org/Publicat/PSR/Briefings/PSR-BriefingApr08.pdf

Chaum, D., Fiat, A., & Naor, M. (1988). Untraceable Electronic Cash. In S. Goldwasser (Ed.), *Advances in Cryptology — CRYPTO' 88* (Vol. 403, pp. 319-327). Santa Barbara, California, USA: Springer New York.

Chirgwin, R. (2011). Bitcoin collapses on malicious trade - Mt Gox scrambling to raise the Titanic. *The Register*. Retrieved September 28, 2011, from http://www.theregister.co.uk/2011/06/19/bitcoin_values_collapse_again/

Clark, J., & Essex, A. (2012). CommitCoin: Carbon Dating Commitments with Bitcoin. *16th International Conference on Financial Cryptography and Data Security*. Bonaire, Caribbean Netherlands.

Dai, W. (1998). b-money. Retrieved from http://www.weidai.com/bmoney.txt

Doherty, S. (2011). All your Bitcoins are ours… | Symantec Connect Community. Retrieved November 10, 2011, from http://www.symantec.com/connect/blogs/all-your-bitcoins-are-ours

Dwork, C., & Naor, M. (1992). Pricing via Processing or Combatting Junk Mail. In E. F. Brickell (Ed.), *Advances in Cryptology — CRYPTO' 92* (Vol. 740, pp. 139-147). Santa Barbara, California, USA: Springer New York.

EIA. (2012). Electric Power Monthly. *U.S. Energy Information Administration*. Retrieved February 22, 2012, from http://www.eia.gov/electricity/monthly/

EURO Kartensysteme GmbH. (2008). Händlerbedingungen - Bedingungen für die Teilnahme am electronic cash-System der deutschen Kreditwirtschaft. Retrieved from http://www.electronic-cash.de/media/pdf/haendlerbedingungen.pdf

Enright, B. (2007). Exposing Stormworm. Retrieved from http://www.scribd.com/doc/2674816/exposing-storm

Europe's Energy Portal. (2010). Retail (end-user) energy prices for households. Retrieved February 22, 2012, from http://www.energy.eu/#domestic

Facebook. (2012). Fact Sheet. Retrieved February 22, 2012, from http://newsroom.fb.com/content/default.aspx?NewsAreaId=22

Franklin, M. K., & Malkhi, D. (1997). Auditable Metering with Lightweight Security. *1th International Conference on Financial Cryptography* (pp. 151-160). Anguilla, British West Indies.

Green500.org. (2012). Green500. Retrieved from http://www.green500.org/

Grigg, I. (2011). Financial Cryptography: BitCoin - the bad news. *financialcryptography.com*. Retrieved February 16, 2012, from http://financialcryptography.com/mt/archives/001327.html

Grinberg, R. (2011). Bitcoin: An Innovative Alternative Digital Currency. *Hastings Science & Technology Law Journal*, *4*, 160-208.

IEA. (2011). *CO2 Emissions from Fuel Combustion 2011*. Retrieved from http://www.iea.org/co2highlights/co2highlights.pdf

IEA. (2012). Electricity/Heat in World in 2009. Retrieved February 15, 2012, from http://www.iea.org/stats/electricitydata.asp?COUNTRY_CODE=29

Kaminsky, D. (2011). Some Thoughts On Bitcoins. *Presentation on slideshare*. Retrieved February 23, 2012, from http://www.slideshare.net/dakami/bitcoin-8776098

Krugman, P. (2011). Golden Cyberfetters. *The New York Times*. Retrieved December 2, 2011, from http://krugman.blogs.nytimes.com/2011/09/07/golden-cyberfetters/

Laurie, B., & Clayton, R. (2004). Proof-of-Work " Proves Not to Work. *3rd Annual Workshop on the Economics of Information Security*. Minnesota, USA. Retrieved from http://www.dtc.umn.edu/weis2004/clayton.pdf

Levy, S. (1994). E-Money (That's What I Want). *Wired Magazine*. Retrieved November 17, 2011, from http://www.wired.com/wired/archive/2.12/emoney.html

Liu, D., & Camp, L. J. (2006). Proof of work can work. *5th Annual Workshop on the Economics of Information Security*. Cambridge, England.

Lübbert, D. (2007). *CO2-Bilanzen verschiedener Energieträger im Vergleich - Zur Klimafreundlichkeit von fossilen Energien, Kernenergie und erneuerbaren Energien*. Retrieved from http://www.bundestag.de/dokumente/analysen/2007/CO2-Bilanzen_verschiedener_Energietraeger_im_Vergleich.pdf

McMillan, R. (2007, October 21). Storm Worm Now Just a Squall. *PCWorld*. Retrieved from http://www.pcworld.com/article/138721/storm_worm_now_just_a_squall.html

Mosenergosbyt. (2012). Electricity tariffs for the population of the city of Moscow in 2012. Retrieved February 22, 2012, from http://www.mosenergosbyt.ru/portal/page/portal/site/personal/tarif/msk

Nakamoto, S. (2008). Bitcoin: A Peer-to-Peer Electronic Cash System. *Network*. Retrieved from https://www.crfdl.org/bitstream/handle/10838/959/bitcoin.pdf?sequence=1

Rajab, M. A., Zarfoss, J., Monrose, F., & Terzis, A. (2007). My botnet is bigger than yours (maybe, better than yours): Why size estimates remain challenging. In N. Provos (Ed.), *1st Workshop on Hot Topics in Understanding Botnets* (p. 5).

Reid, F., & Harrigan, M. (2011). An Analysis of Anonymity in the Bitcoin System. Retrieved from http://arxiv.org/abs/1107.4524

Scròfina, S. (2011). How can Bitcoin be hacked? *Quora.com*. Retrieved February 14, 2012, from https://www.quora.com/How-can-Bitcoin-be-hacked

TOP500.Org. (2012). Top500. Retrieved February 20, 2012, from http://top500.org/

US District Court for the District of Columbia. (2006). *G-Gold indictment* (pp. 1-28). Retrieved from http://www.justice.gov/criminal/ceos/Press Releases/DC egold indictment.pdf

Weiner, S., & Wrigth, J. (2005). Interchange fees in various countries: Developments and determinants. *Review of Network Economics*, *4*(4), 290-323.

Wikipedia. (2012a). Electricity pricing. Retrieved February 22, 2012, a from http://en.wikipedia.org/wiki/Electricity_pricing#Global_electricity_price_comparison

Wikipedia. (2012b). Botnet. Retrieved February 22, 2012, b from http://en.wikipedia.org/wiki/Botnet#cite_note-19

Wikipedia. (2012c). Timeline of Occupy Wall Street. Retrieved February 22, 2012, c from http://en.wikipedia.org/wiki/Timeline_of_Occupy_Wall_Street

Wikipedia. (2012d). 15 October 2011 global protests. Retrieved February 22, 2012, d from http://en.wikipedia.org/wiki/15_October_2011_global_protests#cite_note-atlantic-10

Wikipedia. (2012e). Protests against the Iraq War. Retrieved February 22, 2012, e from http://en.wikipedia.org/wiki/Protests_against_the_Iraq_War

Wikipedia. (2012f). February 15, 2003 anti-war protest. Retrieved February 22, 2012, f from http://en.wikipedia.org/wiki/February_15,_2003_anti-war_protest

Yang, J. (2009). China to Encourage Solar Use. *Wall Street Journal*. Retrieved February 22, 2012, from http://webcache.googleusercontent.com/search?q=cache:deiCSQwS69cJ:online.wsj.com/article/SB124397202782578277.html+China+to+Encourage+Solar+Use&cd=1&hl=de&ct=clnk&gl=de

Zetter, K. (2009). Bullion and Bandits: The Improbable Rise and Fall of E-Gold. *Wired*. Retrieved February 22, 2012, from http://www.wired.com/threatlevel/2009/06/e-gold/

Zhu, Z., Lu, G., Chen, Y., Fu, Z. J., Roberts, P., & Han, K. (2008). Botnet Research Survey. *Computer Software and Applications, 2008 (COMPSAC '08)* (pp. 967-972). Turku, Finland.