

# Feather forking as a positive force: incentivising green energy production in a blockchain-based smart grid

Antonio Magnani\*  
University of Bologna  
Dept. of Computer Science and  
Engineering  
Cesena, Italy  
antonio.magnani@unibo.it

Luca Calderoni  
University of Bologna  
Dept. of Computer Science and  
Engineering  
Cesena, Italy  
luca.calderoni@unibo.it

Paolo Palmieri  
University College Cork  
Dept. of Computer Science  
Cork, Ireland  
p.palmieri@cs.ucc.ie

## ABSTRACT

Climate change represents a serious threat to the health of our planet and imposed a discussion upon energy waste and production. In this paper we propose a smart grid architecture relying on blockchain technology aimed at discouraging the production and distribution of non-renewable energy as the one derived from fossil fuel. Our model relies on a reverse application of a recently introduced attack to the blockchain based on chain forking. Our system involves both a central authority and a number of distributed peers representing the stakeholders of the energy grid. This system preserves those advantages derived from the blockchain and it also address some limitations such as energy waste for mining operations. In addition, the reverse attack we rely on allows to mitigate the behavior of a classic blockchain, which is intrinsically self-regulated, and to trigger a sort of ethical action which penalizes non-renewable energy producers. Blacklisted stakeholders will be induced to provide their transaction with higher fees in order to preserve the selling rate.

## CCS CONCEPTS

• **Hardware** → **Smart grid**; • **Security and privacy** → *Economics of security and privacy*; • **Computer systems organization** → *Peer-to-peer architectures*;

## KEYWORDS

BlockChain, Smart grid, Feather forking

## ACM Reference Format:

Antonio Magnani, Luca Calderoni, and Paolo Palmieri. 2018. Feather forking as a positive force: incentivising green energy production in a blockchain-based smart grid. In *Proceedings of MobiSys'18*. ACM, New York, NY, USA, 6 pages. <https://doi.org/http://dx.doi.org/XXXX.XXXX>

\*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*MobiSys'18, June 10–15, 2018, Munich, Germany*  
© 2018 Association for Computing Machinery.  
ACM ISBN 978-1-4503-5720-3...\$15.00  
<https://doi.org/http://dx.doi.org/XXXX.XXXX>

## 1 INTRODUCTION

Since its introduction [7], blockchain has been one of the most disruptive and promising new technologies, with applications across a wide range of scenarios, well beyond information and communication technology, including economics, law and so forth.

A blockchain is normally defined as a shared distributed ledger of transactions. In general, blockchain technologies offer a distributed platform where participants can record their interactions (such as the exchange or transmission of currency, as in the case of Bitcoin). The transactions composing the blockchain are typically ordered in structured blocks through a linked list, which uses a hash pointer to the preceding block. Hash pointers prevent changes to information stored on previous blocks of the blockchain.

The validity of the recorded transactions is ensured through the shared storage of the ledger information on all the nodes participating to the peer-to-peer blockchain network, as well as a distributed consensus mechanism. The latter provides the way through which new blocks (and the transactions recorded therein) are added to the blockchain. This is one of the areas where research has been the most active. The original Bitcoin proposal is based on a *proof of work* (PoW): in particular, in order to introduce a new block participants need to prove they used a specific amount of computational resources. This is achieved by way of solving a specific computational challenge. Peers participating in the resolution of the challenge are called *miners*. The challenge should be computationally difficult to solve, but simple to verify. Miners collaborate in the distributed effort to solve the challenge, and the probability a specific miner will be the one eventually finding the correct solution should only depend on its relative share of computational resources. Once found by a miner, the solution is shared across the network, which verifies its validity. In the case of Bitcoin, the proof of work is

$$H(\textit{nonce}||P\_hash||Tx_1||\dots||Tx_n) < \textit{target} \quad (1)$$

where  $P\_hash$  is the hash of the preceding block (the hash pointer),  $Tx$  are the  $n$  transactions to be recorded in the current block,  $\textit{target}$  is the challenge difficulty, and  $\textit{nonce}$  is any value for which the hash of the concatenation of the nonce itself, the preceding hash and the transactions is less than the target value. The PoW is therefore the search of a nonce satisfying the equation, normally performed as an exhaustive search (brute force). An advantage of the proof of work approach is embedded resistance to Sybil attacks and double spending [1]. However, it also implies a significant amount of computation (and the associated energy consumption) is wasted for every block created, as the solution will be found by one peer

and all other peers will have worked in vain. PoW blockchains are also prone to frequent forking, where two or more blocks are produced with the same hash pointer. This can happen, for instance, if two peers find the same solution (or two valid solutions) at the same time. Forks are normally resolved by adoption by the network through consensus of the longest fork path, for which the highest amount of computational work has been performed.

An alternative to proof of work is *proof of stake* (PoS). In this common approach, the creation of a new block is assigned probabilistically to one of the network participants according to its relative *stake* (or share) in the network. This reduces significantly the computational resources needed, and the related energy consumption.

In order to motivate peers to participate in the *mining* (the computation needed to solve the challenge), blockchains use an incentive mechanism. In Bitcoin, the peer that succeeds in creating a new block is rewarded by the creation of a new coin which the peer will own. *Transaction fees* can also be introduced to incentivise miners to include specific transactions in the blocks they generate. The fees are normally paid by the peer proposing the transaction once this has been included in a new block.

## 1.1 Contribution

In this paper, we present a distributed blockchain-based architecture for the accounting of energy production and distribution in a smart grid. In particular, we address the incentivisation of green energy production by introducing an ethical mechanism inside the blockchain. Through a derivation of the punitive forking blockchain attack [8], peers in the blockchain can discourage the production of non-renewable energy under certain circumstances, by increasing the related transaction costs.

In this paper, we formalize the feather forking attack and we discuss how it can be applied in the smart grid context for the proposed purpose. We analyse advantages and disadvantages of two well-known models designed to achieve distributed consensus within the blockchain. Finally, we define the architectural aspects of the proposed solution.

## 1.2 Related works

The adoption of blockchain technologies in the smart grid domain has been proposed a number of times. In [5], Li et al. discuss a peer to peer energy trading system relying on blockchain. This model is coupled with a credit-based payment scheme to support fast and frequent energy trading and does not need any trusted intermediary.

Blockchain is not intended to be only adopted in a world-based network (as per the Bitcoin case) but is also suitable for urban contexts [9] and small decentralized markets. In [6] this technology is at the basis of a local energy market between 100 residential households. Local energy trading is performed without the need of a central intermediary.

Many problems connected to the energy Internet may be solved or mitigated through the adoption of blockchain. In [3] the authors discuss how to get rid of the huge amount of money and resources required to purchase and manage energy storage equipment involving the energy Internet. Again, Pop et al. [10] suggest blockchain as

an enabling technology to allow independent energy trading while keeping control of local and area-related consumption through smart meters. This latter system relies on Ethereum platform [12] and may be conveniently used to enhance energy demand and production matching.

SolarCoin [11], a global rewards program for solar electricity generation, should also be mentioned as a successful project relying on blockchain and applied to the energy production domain. The SolarCoin Foundation rewards solar energy producers with blockchain-based digital tokens at the rate of 1 SolarCoin (SLR) per 1 MWh of solar energy produced.

Blockchain technologies also present some drawback. One of the most engaging problem, especially when we deal with smart grids and power supply, is the disproportionate power consumption needed for the proof of work mechanism. Bitcoin represents an evident example of this paradox [2]. As bitcoin and, more in general, blockchain solutions rely on fully distributed systems without any trusted authority, they are vulnerable to attacks by malicious nodes in the network. As discussed in [1], several attacks addressing blockchain technologies were proposed in literature, some of which are based on chain forking. In this paper, we focus on a specific kind of forking attack known as *feather forking* [8]. The main aim of our model consists in a reverse usage of such attack designed to discourage some specific transactions within smart grid domain.

## 2 DESIGN AND DISCUSSION

The collaborative architecture we design is built over a network connecting energy producers and distributors, as well as a national energy authority or regulator. The network follows a standard decentralised peer-to-peer structure, where persistent connections are established between peers. The nature of the network allows peers to enter or exit the network dynamically. The peers (or nodes) composing the network are servers corresponding to energy production facilities, or commercial energy distributors. As such, only nodes recognised by the energy authority participate to the network. In practice, the energy authority could actually provide dedicated hardware to the participants: we call therefore each node a *black box*, as the stakeholders are not allowed to tamper with the functioning of the server. In this context, we assume the parties participating to the network will not launch malicious attacks aimed at disrupting the network or other peers. We assume, however, that the actors will try to minimise their transaction costs. Transaction fees, in particular, will be imposed on specific energy production transactions by non-renewable energy plants, under certain circumstances. As black boxes are assigned to peers by the energy authority, each peer maintains a limited degree of anonymity with respect to other peers [14]. The authority also participates to the network directly, through a supernode of capabilities that are higher than regular black box nodes. The transactions that are recorded by the blockchain architecture are linked to energy production and distribution: producers will record the amount of energy input to the grid, while distributors will record energy collected from the grid, which they will then resell to their customers (households and businesses). The reselling by distributors is not included in the proposed architecture: only transaction to and from the energy grid by energy stakeholders are comprised, excluding final users.

Our architecture does not introduce any currency. In the proposed design, all peers who participate to the network also contribute to the formation of the blockchain, and are as such *miners*.

Wüst and Gervais, in the paper “Do you need a Blockchain?” [13], identify the cases in which a blockchain is a potential solution, and when a problem can instead be solved by traditional means. In the scenario proposed in this paper, the presence of well-known and trusted nodes - as all are certified by a central authority - does not make the adoption of a blockchain-based solution inappropriate. Indeed, despite these characteristics, the proposed setting includes the participation of proactive nodes (individual or as a collective) that can carry out ethical actions independently of the policies suggested by the central authority. The scheme is not aimed at the mere application of energy policies imposed by an authority, but at finding a solution that allows participants to have decision-making autonomy and in which they can organize themselves to counteract or participate in the current policies suggested by the authority (e.g., possible environmental associations may form “green-cooperatives” by creating real pools of votes). In this context, producers of energy from non-renewable sources will also be able to decide whether to ignore or act accordingly to certain policies. A distributed consensus mechanism is therefore necessary, and a solution adopting a centralised database is therefore not feasible. This becomes even more evident with the extension of the energy grid and the consequent inclusion of different central authorities, each with its own energy policies (e.g regional, national and European authorities). In this case, the authorities could make different choices or even act in competition with each other. In such a scenario, the real difference would be made by the individual nodes and their choices: to encourage the production of green energy - preferring it when possible - or to seek a profit through fees?

## 2.1 Proof of Work vs Proof of Stake

The wide applicability of blockchain to the smart grid domain, and in particular to the described scenario, suggests solutions that are not based on the *proof of work* (PoW) model. As discussed in [2], as the number of transactions and the related hash rate increase, so does the computation required and therefore the energy consumed by the peers, which would appear contradictory to an energy efficiency goal. An alternative approach is the *proof of stake* (PoS) model, as such a choice would reduce energy consumption by the peers. However, PoS introduces a number of drawbacks: in particular the *nothing-at-stake* issue [4], where miners without any significant stake and therefore nothing to lose have an incentive in initiating multiple forks. This would allow them to maximize the transaction fee benefits, as validators could generate conflicting blocks on multiple forks with nothing at stake.

Moreover, as in the proposed architecture no currency is generated and transactions record the input or output of energy in the smart grid, the stake could only be based on the volume of energy produced/distributed. This would be imbalanced in a setting where renewable energy (such as solar or wind) producers are normally in higher number but lower capacity with respect to thermonuclear or fossil fuel plants.

For these reasons, rather than adopting the proof of stake model, we opt for a proof of work system where we reduce the hash rate

and therefore the power consumption. Assuming the hardware is provided by an energy authority, this would prevent a race for computational power and therefore an inflation of power consumption. As all peers are restricted to the same hashing power, the architecture introduces a fairness element in the collective behaviour. Adoption of a PoW system also allows for a more direct introduction of the feather forking mechanism we present in the following section, although this would be possible in a PoS scenario as well.

## 2.2 Feather forking

Feather forking is a subtle modification of the more well-known *punitive forking* attack. Punitive forking [8] consists in excluding someone from the blockchain through a systematic and unbounded forking operation with respect to those blocks which contain transactions originating from the blacklisted people. Although this attack is very dangerous (as it could compromise the blockchain usability completely), it is hard to carry out when the attacker does not hold the majority of the hash power of the whole system.

Feather forking is much more affordable to be carried out. This attack can be indeed achieved without the majority of the hash power. Feather forking shares basic concepts with punitive forking but differs from it for a crucial detail: when the attacker announces he will refuse to mine transactions involving a certain person, he also states he is going to fork the chain for a limited number of blocks. Specifically, let us suppose a blacklisted transaction has been inserted in a valid block in the main chain. The attacker could announce he is going to fork the chain starting from the previous block (in order to cut off the new undesired block) and he will keep forking until  $k$  blocks will be added to the main chain after the blacklisted one. As an example, let us suppose  $k = 1$ : a feather forking attack would succeed if the attacker solves two consecutive blocks while other peers do not attach any consequent block to the blacklisted one. Conversely, when a single confirmation is provided for the undesired block (i.e. when a new valid block is added to the main chain after the blacklisted one), the attack fails. This situation is depicted in Figure 1.




Figure 1: Feather forking attack with  $k = 1$ .

A lightweight probabilistic model may be defined in order to understand the impact of such an attack on the network. Let  $\alpha$  be the amount of hash power held by the attacker, where  $\alpha \in [0, 1]$ .

Hash power percentage	$\alpha$	$P(\text{FF}_1)$	Success probability
5%	0.05	0.0025	0.25%
10%	0.1	0.01	1.00%
15%	0.15	0.0225	2.25%
20%	0.2	0.04	4.00%
25%	0.25	0.0625	6.25%
30%	0.3	0.09	9.00%

**Table 1: Feather forking success probability for some reference values of  $\alpha$ . The bounding parameter  $k$  is fixed to 1.**

Assuming all other peers keeps working on the main chain, there is a probability  $\alpha$  that the attacker finds a valid block for the forked chain before someone else does on the main chain. However, we know that the attacker needs to find  $k + 1$  consecutive blocks in order for the attack to succeed.

*Definition 2.1.* Given a bound  $k$ , we define a *successful feather forking attack*, and we refer to it as  $\text{FF}_k$ , the event when an attacker connects  $k + 1$  consecutive blocks to a forked chain before other peers connect at least  $k$  consecutive blocks to the main chain.


For instance, when  $k = 1$ , this condition is reached when the attacker finds two blocks while the rest of peers find none. Thus, the probability for  $\text{FF}_1$  to occur is:

$$P(\text{FF}_1) = \alpha^2. \quad (2)$$

Let us assume  $\alpha = 0.1$  (i.e. the attacker holds 10% of the hash power). When he performs feather forking he succeeds on the average one time in a hundred. Again,  $\alpha = 0.2$  implies a 4% of failures when someone attempts to mine a block containing a blacklisted transaction, and so forth. Some reference values are reported in Table 1.

It is important to mention here that this attack can drive several peers to follow the same behaviour, as blocks involving blacklisted transactions have a certain probability to be excluded from the chain (while other blocks have not). Hence, the attacker may trigger a bigger impediment than expected. In practice, however, the probability of a successful feather fork is always marginal: the best line of defence for impacted peers is to pay a higher transaction fee, to ensure the block containing their transactions is more valuable to the community of miners, and therefore it will be included in the main chain. Therefore, the main objective of feather forking is to increase the costs for the victims of the attack, rather than actually succeeding in forking the chain.

The main idea behind the proposed model is to change perspective and use this attack as a positive force: in particular, the energy authority may carry out feather forking attacks to discourage non-renewable energy production when this is not needed given the current consumption level. Depending on the hash power the authority has on the supernode it controls with respect to the other peers, targeted stakeholders will pay a higher fee in order for their transactions to be validated.



**Figure 2: The proposed system architecture, where energy is distributed from the producers to the grid, managed by the central energy authority, and then onward to the commercial energy distributors.**

### 2.3 System architecture

The distributed system architecture is comprised of the black box nodes, each related to an energy producer or distributor, and the energy authority or regulator, which controls one or more supernodes with a significant proportion of the hash power. In the following, we refer to the latter entity as the *Energy Central Authority (ECA)*. Two other classes of stakeholders compose the network of peers: *Energy Producers (EP)*, which are energy production facilities such as a wind farm or a gas power plant; and *Distributors (D)*, which buy energy from the grid and resell it to individual businesses or households. We note that even in the case two energy production plants are owned by the same company, each plant would still be an independent node on the network, with an individual black box.

Through a standard peer-to-peer structure, each ordinary (black box) node participating to the network will have access to other nodes, but will maintain a limited number of active connections. The number of connections in Bitcoin, 8, seems suitable for this scenario as well. Peers will also have a maximum number of incoming connections (e.g. 125). Among the connections maintained, peers should however always select an ECA supernode. Connections can be kept alive through a system similar to that used in Bitcoin: an Hello Message to keep the connection alive, and discard the peer if nothing is heard in the last 30 minutes. The architecture, comprised of stakeholders divided in 3 categories, is depicted in Figure 2.

We imagine the ECA role will be played by a national or international energy authority or regulator, which we assume would have the legal mandate to impose costs on non-renewable energy, as in the proposed objective. The ECA will act on the basis of the current circumstances of the network, by partially regulating the market: by increasing the costs for specific energy sources, the market can be expected to self-regulate in the direction of increased usage of renewable energy. As the decision as whether or not to impose such costs will ultimately lie on the ECA, its application can depend on the network circumstances: for instance, we can imagine a scenario where energy consumption is greater than energy production, for

which the ECA may decide to temporarily suspend the mechanism. The ECA is also responsible for the production of black boxes, which are assigned to each other participant in the network. Where the black box is assigned to an energy production facility EP, this will monitor the energy that is input to the electrical grid; where the black box is instead assigned to an energy distributor D (such as a wholesale customer that resells to businesses and households), it will monitor the energy output or consumed. The transaction recorded in the blockchain will track these inputs and outputs of energy. Each black box has the same computational resources, and therefore will participate in the blockchain with an identical hash power. The supernode held by the ECA has instead a variable hash power, aimed at maintaining a specific percentage of the overall network hash rate. All nodes will have a public and private key pair, and also maintain an up-to-date copy of the blockchain locally. Only the ECA, however, has direct knowledge of the identity of all nodes. This will allow it to impose costs through the feather forking mechanism to specific EPs, which it knows use a certain kind of energy source.

## 2.4 Transactions

In this section, we discuss how the ECA can impose a cost on transactions by specific EP peers through the feather forking mechanism, when the EP uses a non-renewable energy source and the ECA conditions are met.

Transactions in the proposed blockchain are similar to those in Bitcoin, but with a significant difference: no coin or currency is produced by the miners. Therefore, the act of mining in itself will not provide any reward to a miner, and the only reward are transaction fees. We assume the black boxes will be designed to mine independently of the reward, as stakeholders have external interests in participating to the network.

As described in Section 2.2, when the ECA decides to impose a cost on EPs using a particular source of energy (e.g. coal), it will announce through the supernode its intention to fork blocks containing transactions of the targeted EPs, following the feather forking strategy and using a declared hash power. This hash rate can be adjusted by the ECA to increase or decrease the cost to the EP, and according to the grid circumstances.

In order to have transactions validated, and therefore being able to input energy to the grid, the targeted EPs will have to add a transaction fee to their transactions. The fee will incentivise peers to include the transaction on the block, and therefore not participate in the fork. The mechanism is in fact successful if the EPs pay a transaction fee, and not necessarily if the fork is adopted as main branch. The increased costs of energy production imposed by the ECA on the EPs is dynamic and adjustable, as the fee will be proportional to the probability of success of the feather fork (see Table 1). The cost can therefore be calculated in advance by the ECA, with the target of a decreased usage of non-renewable energy sources, without blocking altogether access to the grid.

In Table 2, we distinguish three types of transactions: EP to ECA, ECA to D, and any miner (peer) to ECA.

In the first case, the transaction will have as sender address the public key of the energy production facility EP, and as receiver the public key of the energy authority ECA. The transaction records

Sender <i>PubKey</i>	Receiver <i>PubKey</i>	Fee	Feather Forking
EP	ECA	✓	✓
ECA	D	✗	✗
Miner	ECA	✗	✗

**Table 2: Transaction classes in the proposed smart grid blockchain.**

the energy the EP inputs to the grid. Each transaction could for instance record that a fixed amount of energy has been input, or the amount relative to a fixed time duration, in order to standardise transactions. When the EP decides to add a transaction fee (e.g. to counter a feather fork), this will be earned by the peer mining the block containing the transaction. Fees are also energy: the EP will therefore have to produce more than it sells to the grid.

Once the blockchain has progressed, and transactions are confirmed (in the way of Bitcoin), the ECA can distribute the energy for the recorded transactions to any wholesale energy distributor D. This interaction is recorded in the second type of transaction.

The third and final transaction class is the withdrawals of earning due to transaction fees by the miners who “earned” energy in this way. The peers that have mined blocks and earned the related fees are able to redeem their value through the ECA. The transaction records the fees being passed to the ECA, which will pay the peers an equivalent amount externally to the system after confirmation, in the same way energy would be paid for by Ds.

As the transactions belonging to the second and third group are not related to a specific energy source, they cannot be subject to blacklisting and forking by the ECA. In the event of a successful fork, normally unlikely but probabilistically possible, the ECA will earn any transaction fee, and the affected EPs will not have recorded the energy they input into the grid. The ECA will therefore discount the same amount against fees it would have imposed on the EPs over time. Transaction fees will be, in this case, earned by the ECA supernode, but the ECA can decide to redistribute them to renewable energy EPs.

Through the transaction fees imposed by the ECA thanks to opportunistic feather forking, the proposed model introduces a cost on non-renewable energy producers, and incentivises green energy production through the distribution of the transaction fees. While transaction fees can be earned by non-renewable energy peers as well, statistically this will only partially reduce their cost. This can also be predicted by the ECA, and it can be factored in its decision on the fee level to be imposed, and the consequent hash rate and maximum number of blocks for the attempted fork.

## 3 CONCLUSION

In this paper we introduced a novel blockchain-based system aimed at the regulation of energy production and distribution. A specific focus was posed on the type of energy which the producer plugs in the grid. Specifically, we discussed a tailored application of the feather forking attack designed to discourage the production of non-renewable energy. This technique seems to be promising and might be adopted to enhance ethical smart grid systems where both

a central authority and each participating peer collaborate for a greener environment.

## REFERENCES

- [1] Mauro Conti, Sandeep Kumar E, Chhagan Lal, and Sushmita Ruj. 2017. A Survey on Security and Privacy Issues of Bitcoin. *CoRR* abs/1706.00916 (2017). arXiv:1706.00916 <http://arxiv.org/abs/1706.00916>
- [2] P. Fairley. 2017. Blockchain world - Feeding the blockchain beast if bitcoin ever does go mainstream, the electricity needed to sustain it will be enormous. *IEEE Spectrum* 54, 10 (October 2017), 36–59. <https://doi.org/10.1109/MSPEC.2017.8048837>
- [3] Tao Fan, Qingsu He, Erbao Nie, and Shaozhen Chen. 2018. A study of pricing and trading model of Blockchain & Big data-based Energy-Internet electricity. *IOP Conference Series: Earth and Environmental Science* 108, 5 (2018), 052083. <http://stacks.iop.org/1755-1315/108/i=5/a=052083>
- [4] Wenting Li, Sébastien Andreina, Jens-Matthias Bohli, and Ghassan Karame. 2017. Securing Proof-of-Stake Blockchain Protocols. In *Data Privacy Management, Cryptocurrencies and Blockchain Technology - ESORICS 2017 International Workshops, DPM 2017 and CBT 2017, Oslo, Norway, September 14-15, 2017, Proceedings (Lecture Notes in Computer Science)*, Joaquín García-Alfaro, Guillermo Navarro-Arribas, Hannes Hartenstein, and Jordi Herrera-Joancomartí (Eds.), Vol. 10436. Springer, 297–315. [https://doi.org/10.1007/978-3-319-67816-0\\_17](https://doi.org/10.1007/978-3-319-67816-0_17)
- [5] Z. Li, J. Kang, R. Yu, D. Ye, Q. Deng, and Y. Zhang. 2017. Consortium Blockchain for Secure Energy Trading in Industrial Internet of Things. *IEEE Transactions on Industrial Informatics* PP, 99 (2017), 1–1. <https://doi.org/10.1109/TII.2017.2786307>
- [6] Esther Mengelkamp, Benedikt Notheisen, Carolin Beer, David Dauer, and Christof Weinhardt. 2018. A blockchain-based smart grid: towards sustainable local energy markets. *Computer Science - R&D* 33, 1-2 (2018), 207–214. <https://doi.org/10.1007/s00450-017-0360-9>
- [7] Satoshi Nakamoto. 2009. Bitcoin: A peer-to-peer electronic cash system.
- [8] Arvind Narayanan, Joseph Bonneau, Edward Felten, Andrew Miller, and Steven Goldfeder. 2016. *Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction*. Princeton University Press, Princeton, NJ, USA.
- [9] Alessandra Pieroni, Noemi Scarpato, Luca Di Nunzio, Francesca Fallucchi, and Mario Raso. 2018. Smarter City: Smart Energy Grid based on Blockchain Technology. *International Journal on Advanced Science, Engineering and Information Technology* 8, 1 (2018), 298–306.
- [10] Claudia Pop, Tudor Cioara, Marcel Antal, Ionut Anghel, Ioan Salomie, and Massimo Bertoini. 2018. Blockchain Based Decentralized Management of Demand Response Programs in Smart Energy Grids. *Sensors* 18, 1 (2018), 162. <https://doi.org/10.3390/s18010162>
- [11] SolarCoin Foundation. 2018. *SolarCoin: A blockchain-based solar energy incentive*. Technical Report. SolarCoin Foundation.
- [12] Gavin Wood. 2017. *Ethereum: A Secure Decentralised Generalised Transaction Ledger*. Technical Report. Ethereum. EIP-150 REVISION.
- [13] Karl Wüst and Arthur Gervais. 2017. Do you need a Blockchain? *Cryptology ePrint Archive*, Report 2017/375. <https://eprint.iacr.org/2017/375>.
- [14] Niels Zeilemaker, Zekeriya Erkin, Paolo Palmieri, and Johan A. Pouwelse. 2013. Building a privacy-preserving semantic overlay for Peer-to-Peer networks. In *2013 IEEE International Workshop on Information Forensics and Security, WIFS 2013*. IEEE, 79–84. <https://doi.org/10.1109/WIFS.2013.6707798>