



UCL

UCL Centre for Blockchain Technologies

Discussion Paper Series

Q3 2021



Discussion Paper

Energy Footprint of Blockchain Consensus Mechanisms Beyond Proof-of-Work

*Moritz Platt,^{1,2} Johannes Sedlmeir,³ Daniel Platt,⁴
Jiahua Xu,¹ Paolo Tasca,¹ Nikhil Vadgama,¹ Juan Ignacio Ibañez¹*

¹Centre for Blockchain Technologies, University College London, London, UK

²Department of Informatics, King's College London, London, UK

³FIM Research Center, University of Bayreuth, Bayreuth, Germany

⁴Department of Mathematics, Imperial College London, London, UK

Abstract

Popular distributed ledger technology (DLT) systems using proof-of-work (PoW) for Sybil attack resistance have extreme energy requirements, drawing stern criticism from academia, businesses, and the media. DLT systems building on alternative consensus mechanisms, foremost proof-of-stake (PoS), aim to address this downside. In this paper, we take a first step towards comparing the energy requirements of such systems to understand whether they achieve this goal equally well. While multiple studies have been undertaken that analyze the energy demands of individual Blockchains, little comparative work has been done. We approach this research question by formalizing a basic consumption model for PoS blockchains. Applying this model to six archetypal blockchains generates three main findings: First, we confirm the concerns around the energy footprint of PoW by showing that Bitcoin's energy consumption exceeds the energy consumption of all PoS-based systems analyzed by at least three orders of magnitude. Second, we illustrate that there are significant differences in energy consumption among the PoS-based systems analyzed, with permissionless systems having an overall larger energy footprint. Third, we point out that the type of hardware that validators use has a considerable impact on whether PoS blockchains' energy consumption is comparable with or considerably larger than that of centralized, non-DLT systems.

Keywords: Blockchain, Carbon Footprint, Distributed Ledger Technology, Proof-of-Stake, Sustainability.

Energy Footprint of Blockchain Consensus Mechanisms Beyond Proof-of-Work

Moritz Platt^{1,2}, Johannes Sedlmeir³, Daniel Platt⁴, Jiahua Xu¹,
Paolo Tasca¹, Nikhil Vadgama¹, Juan Ignacio Ibañez¹

¹Centre for Blockchain Technologies, University College London, London, UK

²Department of Informatics, King’s College London, London, UK

³FIM Research Center, University of Bayreuth, Bayreuth, Germany

⁴Department of Mathematics, Imperial College London, London, UK

moritz.platt@kcl.ac.uk, {jiahua.xu, p.tasca, nikhil.vadgama, j.ibanez}@ucl.ac.uk

johannes.sedlmeir@fim-rc.de, daniel.platt.17@ucl.ac.uk

Abstract—Popular permissionless distributed ledger technology (DLT) systems using proof-of-work (PoW) for Sybil attack resistance have extreme energy requirements, drawing stern criticism from academia, business, and the media. DLT systems building on alternative consensus mechanisms, foremost proof-of-stake (PoS), aim to address this downside. In this paper, we take a first step towards comparing the energy requirements of such systems to understand whether they achieve this goal equally well. While multiple studies have been undertaken that analyze the energy demands of individual blockchains, little comparative work has been done. We approach this research gap by formalizing a basic consumption model for PoS blockchains. Applying this model to six archetypal blockchains generates three main findings: First, we confirm the concerns around the energy footprint of PoW by showing that Bitcoin’s energy consumption exceeds the energy consumption of all PoS-based systems analyzed by at least two orders of magnitude. Second, we illustrate that there are significant differences in energy consumption among the PoS-based systems analyzed, with permissionless systems having an overall larger energy footprint. Third, we point out that the type of hardware that validators use has a considerable impact on whether PoS blockchains’ energy consumption is comparable with or considerably larger than that of centralized, non-DLT systems.

Index Terms—Blockchain, Carbon Footprint, Distributed Ledger Technology, Proof-of-Stake, Sustainability

I. INTRODUCTION

In distributed ledger technology (DLT) systems, consensus mechanisms fulfill multiple purposes surrounding the proposal, validation, propagation, and finalization of data [1]. A critical problem for DLT systems are Sybil attacks in which an attacker creates an artificially large number of bogus identities [2] to skew the results of majority decisions on the admission and order of transactions. In permissioned networks, gatekeeping strategies can be applied that limit access to a network to previously vetted actors [3], thereby preventing such attacks.

However, for permissionless networks, in which participants can participate in consensus without any control [4], more complex mechanisms need to be applied to combat Sybil attacks. Proof-of-work (PoW) is an example of a Sybil attack resistance scheme that has been used in most early cryptocurrencies such as Bitcoin [5]. To counteract Sybil attacks, PoW uses cryptographic puzzles of configurable difficulty with efficient

verification such that it becomes computationally expensive for attackers to interfere with consensus [6]. However, by this design, the energy consumption of a PoW-based cryptocurrency strongly correlates with its market capitalization, leading to an extreme energy demand for popular implementations [7]. For instance, the electricity demand of Bitcoin is now in the same range as that of entire industrialized nations [8]. Against this backdrop, many alternatives to PoW have been proposed that do not rely on extensive computational effort [9]. Among those is proof-of-stake (PoS) in which participants with larger holdings of a cryptocurrency have larger influence in transaction validation. While PoS is generally understood as being more energy efficient than PoW, the exact energy consumption characteristics of PoS-based systems, and the influence that network throughput has on them, are not widely understood.

Two main approaches to quantify the energy consumption of a DLT system have been assumed in the past. One is to measure the consumption of a representative participant node and then extrapolate from this measurement. An alternative approach is to develop a mathematical model that includes core metrics of a DLT system to calculate its energy consumption. Extensive research efforts have cumulated in best practices for determining the energy consumption of DLT systems [10]. So far, most work has focused on PoW blockchains,¹ and some research has investigated individual non-PoW systems. In this paper we propose a simple energy consumption model, applicable to a broad range of DLT systems that use PoS for Sybil attack resistance. Specifically, this model considers the number of validator nodes, their energy consumption, and the network throughput based on which the energy consumption per transaction is estimated. We present the results of applying this model to six PoS-based systems. Our results illustrate that, while negligible compared to PoW, the energy consumption of PoS systems can still vary significantly.

The next section surveys related work in both experimental and mathematical models. We then review selected PoS systems

¹While not all distributed ledger technologies organize their data into chains of hash-linked blocks, the term “blockchain” is customarily interchangeable with DLT. In the remainder, we follow this convention for simplicity.

– Ethereum 2.0, Algorand, Cardano, Polkadot, Tezos, and Hedera – and describe their relevant architectural features. In the following section, we introduce our model in more detail and describe how the underlying data was obtained. We apply the model to the systems selected, present the comparative results, and discuss limitations. Finally, we conclude our study with potential avenues for future research.

II. RELATED WORK

We conducted an informal literature review using the search string ("Blockchain" OR "DLT" OR "Distributed Ledger") AND ("Energy Consumption" OR "Energy Demand" OR "Electricity Demand" OR "Carbon Footprint") year:[2008 TO *] on the Bielefeld Academic Search Engine (BASE). We thereby obtained 413 results of various prior work on analyzing the energy demand of different DLT systems, with a significant focus on PoW blockchains in general and specifically Bitcoin. Commonly, models take one of the following two forms.

Experimental Models: The first form revolves around conducting experiments using mining hardware and measuring its actual energy consumption, as done by Igumenov *et al.* with different configurations of computational resources [11]. This approach has been used to derive consumption characteristics for different usage scenarios. The “BCTMark” framework [12], for instance, allows for the deployment of an entire experiment stack, including the DLT system under test. Using load generators, a realistic network workload can be created. The effects on the energy consumption of this setup under varying loads can subsequently be measured via energy sensors connected to the testbed. An experimental study on the energy consumption of the non-PoW XRP ledger demonstrates that customizing validator hardware can yield reductions in energy demand [13]. Metrics reported for common cryptocurrencies have been combined with testbed experiments to model the energy consumption behaviors of various consensus algorithms [14].

Mathematical Models: An alternative method is to quantify assumptions about the environment in which a DLT system operates. Often, such models use a “top-down” approach that relies on publicly observable factors – such as hash rate in the case of Bitcoin – and associate them with common mining hardware or even seek to determine the hardware used via surveys [10]. The papers of Gellersdörfer *et al.* [15], Küfeoglu and Özkuran [16], and Zade *et al.* [17] are examples of this hash rate-based approach. Sedlmeir *et al.* [7] undertake a basic comparison of different DLT architectures with the conclusion that the energy consumption differs significantly depending on the design chosen. A further study by the same authors [18] refines previous models for Bitcoin’s power consumption, such as the one by Vranken [19], and emphasizes that the driving forces behind power consumption are the Bitcoin price and the availability of cheap electricity. Eshani *et al.* [20] use a linear regression model to predict Ethereum’s energy consumption based on the observed hash rate and difficulty level; however, the use of simplistic interpolation techniques alone is likely not an appropriate method for PoW

blockchains [10]. Powell *et al.* [21] derive a mathematical model for the energy consumption of the PoS-based Polkadot blockchain by extrapolating from the power demand of a single validator machine.

III. SYSTEMS REVIEWED

Our comparison set includes DLT systems with high market capitalization that share a critical common denominator: using a PoS-based consensus algorithm. In PoS, validators with a higher stake – often in the form of the DLT system’s native currency – influence the transaction validation more. Thus, the scarce resource of energy to avoid Sybil attacks in PoW is replaced by the scarce resource of capital in the cryptocurrency [7]. Despite the commonalities, these systems differ in a range of other aspects, such as the minimum thresholds to validate and delegate, the necessity to lock-up tokens in order to stake (“bonding”), and the architecture of incentives consisting of penalties (“slashing”) and rewards beyond transaction fees (“block rewards”) (cf. Table II).

When it comes to energy consumption, however, differences in the accounting model, transaction validation mechanism, and node permissioning setting (cf. Table I), together with the architectural design of each system’s specific PoS protocol, are of particular relevance. In this section, we describe each of the PoS-based systems with a focus on those aspects. A full exploration of all possible factors is beyond the scope of this paper.

Ethereum 2.0: Ethereum is a highly popular permissionless blockchain that is currently transitioning from PoW (Ethereum 1.0) to PoS (Ethereum 2.0). In Ethereum 1.0, every full node needs to store all 350 GB of current state data.² However, the storage of the full history of all transactions is used by archive nodes only. There are also light nodes storing only the header chains and requesting everything else from a full node on which they depend.³ The sharding proposal (Ethereum 2.0 Phase 1), designed to limit compute, storage, and bandwidth needs, is not yet active.

Algorand: Algorand is a permissionless, account-based system where relay nodes store the entire ledger and non-relay nodes store approximately 1,000 blocks. A proposal to limit storage needs through transaction expiration and sharding (“Vault”) is not yet active [22].

Cardano: Cardano is also permissionless and the only unspent transaction output (UTXO)-based system in our comparison set. In Cardano, nodes store all transactions ever made. Its proposal for sidechains and sharding (“Basho”) is not yet active. There is a probability of being selected as block-proposer for an epoch weighted by stake. However, it is possible to delegate the stake to a stake pool, whose manager receives rewards when the pool is selected and then shares them with the delegators. Rewards are diminishing with the pool size if a pool is so large that it exceeds a saturation parameter. Non-selected stakers verify proposed blocks [23].

²<https://ethereum.org/en/developers/docs/storage/>

³<https://ethereum.org/en/developers/docs/nodes-and-clients/>

Polkadot: In Polkadot’s permissionless nominated proof-of-stake (NPoS), each node can delegate stake to up to 16 validators, among which the stake is always divided equally. Rewards to validators are proportionate to validation work, not to their stake. Polkadot also distinguishes between archive nodes (storing all past blocks), full nodes (256 blocks), and light nodes (storing only runtime and current state, but no past blocks). The first five shards (“parachains”) have been already auctioned on the testnet “Kusama” but have not been deployed in the main chain.

Tezos: In Tezos’ permissionless liquid proof-of-stake (LPoS), stake can also be delegated. Some delegates are block producers, other delegates verify; both receive rewards for it proportional to their stake [24]. Nodes have a “full mode” storing the necessary data needed to reconstruct the complete ledger state since the genesis block, but not contextual data from a checkpoint onwards; an “archive mode” where all blockchain data since the genesis block including contextual data such as past balances or staking rights beyond the checkpoints are stored; and “rolling mode” that only stores the minimal data that is necessary to validate blocks.

Hedera: In contrast to the other five systems studied, Hedera is permissioned network that uses a directed acyclic graph (DAG)-based data structure to store the transaction history (cf. Table I) and applies PoS [25]. The network has its consensus nodes run solely by its council members at the moment, with the plan to open up to permissionless nodes in the future⁴. Transactions do not form blocks, but are spread through a “gossip about gossip” protocol where new information obtained by any node is spread exponentially fast though the network [26]. The consensus calculation takes the form of a weighted average of all gossiping nodes’ information such as transaction order, with the weight proportionate to a node’s stake.

Platform	Accounting Model		Data structure		Permissioning	
	Account	UTXO	Block	DAG	P'ned	P'less
Ethereum 2.0	•		•			•
Algorand	•		•			•
Cardano		•	•			•
Polkadot	•		•			•
Tezos	•		•			•
Hedera	•			•	•	

Table I

COMPARISON OF THE ANALYZED DLT SYSTEMS IN ACCOUNTING MODEL, DATA STRUCTURE, AND NODE PERMISSIONING SETTING

IV. METHOD

Our model differs from previous work (cf. Section II) in that we focus on the energy consumption per transaction, as opposed to the overall energy consumption of an entire system. Nevertheless, existing models can be combined with additional data arising from the scientific literature, reports, and public ledger information to form a baseline that can be used

⁴<https://help.hedera.com/hc/en-us/articles/360000674017-Is-the-Hedera-public-network-permissioned-or-permissionless->

Platform	Bonding	Slashing	Rewards
Ethereum 2.0	Yes	Yes	Yes
Algorand	No	No	Yes
Cardano	No	No	Yes
Polkadot	Yes	Yes*	Yes
Tezos	No	Yes	Yes
Hedera	No	No	No

Table II

PROPERTIES OF THE PoS PROTOCOLS USED FOR SYBIL ATTACK RESISTANCE IN THE DLT SYSTEMS ANALYZED

to avoid time-consuming experimental validation. Powell *et al.* [21] define an elementary mathematical model for the energy consumption of the Polkadot blockchain that can be generalized as

$$p_t = p \cdot n_{\text{val}}, \quad (1)$$

where p_t is the overall average power the DLT system consumes, p is average power consumed by a validator node, and n_{val} is the number of validator nodes. Due to the comparatively low computational effort associated with PoS and the intentionally relatively low throughput of permissionless blockchains to avoid centralization because of compute, bandwidth, or storage constraints [27], it can be assumed that validating nodes run on similar types of commodity server hardware, irrespective of the network load.

Under this assumption, the overall energy need of such a protocol is solely contingent on the number and hardware configuration of validator nodes. In the context of this paper, we only consider the energy footprint of the consensus mechanisms itself. We, therefore, only consider *validators*⁵, i.e., nodes that actively participate in a network’s consensus mechanism by submitting and verifying the proofs necessary for Sybil attack resistance [1]. The overall number of nodes, including other *full nodes* that replicate the transaction history without participating in consensus, is likely higher for all systems analyzed. A key model parameter, therefore, is the number of validator machines running concurrently (n_{val}). This number can be established reliably, since it is stored on-chain as a key aspect of any PoS-based protocol. Table III shows the number of validators currently operating on each of the networks considered.

Platform	# Validators	TPS Cont. (tx/s)	TPS Max. (tx/s)
Ethereum 2.0	183 753		3000
Algorand	1126	9.85	1000
Cardano	2958	0.36	257
Polkadot	297	0.12	1000
Tezos	399	1.70	40
Hedera	21	48.20	10 000

Table III

THE CURRENT NUMBER OF VALIDATORS, CONTEMPORARY THROUGHPUT, AND THE UPPER BOUND OF THROUGHPUT POSTULATED (CF. APPENDIX A).

⁵A node fulfilling this role goes by various names, e.g., “stake pool” for Cardano, or “baker” for Tezos.

Energy consumption per transaction: To arrive at an energy consumption per transaction metric (c_{tx}), the number of transactions per unit of time needs to be considered. The actual numbers are dynamic and fluctuate over time. The contemporary network throughput (*Cont.*) is defined as the actual throughput a system experienced recently. As a key metric, this can be derived from approximate timestamps that are associated with transactions on public ledgers. The maximum postulated sustainable system throughput (*Max.*) of the different protocols is derived from casual sources (cf. Appendix B). Note that these postulated figures are likely optimistic, that means, not necessarily reliable, as they originate not from controlled experiments but are anecdotal or come from promotional materials. However, we consider these estimates acceptable as they have no direct influence on the energy consumption per transaction for a fixed contemporary network throughput. They merely dictate the domain of the consumption function $f_{c_{\text{tx}}}(l)$ that calculates the consumption per transaction depending on the overall system throughput l (measured in tx/s). Treating the average power consumed by a validator node (p , measured in W) as a constant means that an inverse relationship between consumption per transaction (c_{tx}) and system throughput (l) can be established within the bounds of $(0, l_{\text{max}}]$:

$$f_{c_{\text{tx}}}(l) = \frac{n_{\text{val}} \cdot p}{l}. \quad (2)$$

Modelling c_{tx} as a function of the number of transactions per second: Equation (2) depends on two variables: n_{val} and l . We will now present a model for c_{tx} that depends on one variable, namely l , only. Data from the Cardano blockchain [28] suggests that the number of validators n_{val} and the number of transactions per second l are positively correlated. Namely, Pearson’s correlation coefficient⁶ for n_{val} and l for 375 data points from 29 Jul 2020 to 7 Aug 2021 is 0.80. The correlation coefficient for n_{val} delayed by 28 days and l (not delayed) for the same data is 0.87. This is plausible for the following reason: as the total number of users in a permissionless system increases, a share of these new users becomes validators and another non-disjoint share executes transactions, meaning that n_{val} and l are positively correlated. For permissioned systems, it is still conceivable that the number of validators and throughput are linearly dependent and positively correlated because as new partner organizations are invited to run validator nodes, these partners may decide to use the system for their own applications, thereby increasing the number of transactions. We also observe that in the case of Hedera, the number of validators and throughput are positively correlated: the number of validator nodes has been continuously increasing; and throughput, while fluctuating from month to month, has increased year-to-year (cf. Appendix A). Furthermore, it can be observed for the Algorand and Hedera blockchains that n_{val} and l have increased from July to August 2021. On the Polkadot blockchain, n_{val} has remained constant from February to July 2021. An exception is the Tezos

⁶The correlation coefficient takes values in $[-1, 1]$ and a value of ± 1 would imply that n_{val} is an affine function in l .

blockchain for which n_{val} has decreased while l has increased from February to August 2021. This trend has so far held true throughout the lifetime of the Tezos blockchain. We note that in this case an affine function is not appropriate to model the dependence of n_{val} on l , because n_{val} would become negative for large values of l . We will still compute the affine best approximation of n_{val} in terms of l for the Tezos blockchain, as it is an approximation of the first Taylor polynomial of n_{val} , and therefore a local model for n_{val} .

For simplicity we assume that the correlation is perfect, i.e., $n_{\text{val}} = \kappa + \lambda \cdot l$ for some $\kappa, \lambda \in \mathbb{R}, \lambda > 0$, and using (2) we obtain

$$f_{c_{\text{tx}}}(l) = \frac{(\kappa + \lambda l) \cdot p}{l}. \quad (3)$$

Because we could not obtain high-resolution historic data for Algorand, Polkadot, Tezos, and Hedera, we will later on compute κ, λ based on two data points. For Cardano, we use linear regression implemented as ordinary least squares regression to compute κ, λ that have the maximum likelihood of modelling $f_{c_{\text{tx}}}(l)$ under the assumption that $f_{c_{\text{tx}}}(l)$ is an affine function with Gaussian noise. The resulting values for κ, λ can be found in Table IV.

Platform	κ	λ
Algorand	102.8	103.9
Cardano	1267.8	2959.2
Polkadot	297.0	0.0
Tezos	440.7	-24.6
Hedera	7.6	0.3

Table IV
ESTIMATES FOR κ, λ FOR DIFFERENT DLT PLATFORMS USED IN (3) TO MODEL THE NUMBER OF VALIDATORS DEPENDING ON THE NUMBER OF TRANSACTIONS PER SECOND.

Hardware Type and Compute Resource Utilization Considerations: In stark contrast to energy-intensive PoW systems, in PoS, the computational effort relating to the participation in the consensus protocol can practically be considered independent of extraneous factors like cryptocurrency capitalization. Numerous factors influence the overall energy consumption of a server with central processing unit (CPU) activity, hard disk drive operations and cooling contributing most significantly to it [29]. Consensus-related energy demand in PoS is generally constant, meaning it occurs irrespective of system load [29]. Energy demand relating to CPU time and input/output operations is, however, highly load-dependent [30]. Therefore, a realistic energy consumption estimate for a validator node needs to factor in both the minimum hardware requirement (i.e., how many CPU cores or what amount of memory is required) as well as the utilization of that hardware.

Since it is close to impossible to determine which type of hardware is used by validators in actuality, we use an approximation derived from industry recommendations. For permissionless systems and permissioned systems dramatically different hardware recommendations are put forward. The permissionless systems analyzed in this study, all traditional blockchains with comparatively large numbers of validators

running full nodes that verify every transaction [27], demand comparatively low-powered hardware. Hedera, the only permissioned system analyzed here, constitutes a high-transactions per second (tps) system. Such systems are characterized by a small number of nodes maintains consensus [27]. As such, the network performance is determined by the lowest-performing validator node⁷. Therefore, in order to achieve the postulated maximum throughput values, highly performant server hardware is demanded by the network operator. We assumed that similar high-tps systems would have energy requirements in the same range. This explains the difference in energy consumption per validator node between Hedera and the other traditional Blockchain systems.

Config.	Hardware Type	Exemplar	Demand (W)
Minimum	Small single-board computer	Raspberry Pi 4	5.5
Medium	General-purpose rackmount server	Dell PowerEdge R730	168.1
Maximum	High-performance server	Hewlett Packard Enterprise ProLiant ML350 Gen10	328

Table V
CONCEIVABLE UPPER AND LOWER BOUNDS FOR THE POWER DEMAND OF A VALIDATOR MACHINE

To capture the uncertainty regarding appropriate hardware and expected hardware utilization in the model, three different validator configurations are considered (cf. Table V): a single-board computer, a general-purpose rackmount server for midsize and large enterprises, and a high-performance server. For all configurations, hardware utilization based on typical workloads is assumed (cf. Appendix C). For traditional blockchains, we assume a power demand in the minimum to medium range (5.5 W to 168.1 W). For high-tps systems, the medium to maximum range (168.1 W to 328 W) is assumed.

V. RESULTS

Table VI illustrates the application of the models described in (2) and (3) to estimate the energy consumption of the protocols considered under contemporary throughput, i.e., based on recent throughput measures (cf. Section IV). To facilitate a broad overview, we also provide the global system-wide consumption of each DLT system according to the model. Furthermore, the table presents two estimates for energy consumption per DLT system: an *optimistic* estimate assuming validator nodes are operated on the lower bound of the system range and a *pessimistic* estimate that assumes validators utilize hardware on the higher bound (cf. Table V). As the merge of Ethereum mainnet with the beacon chain is outstanding, no contemporary throughput figures for Ethereum 2.0 can be established. Instead, a broad projection between a lower bound, the current throughput of the Ethereum blockchain (15.40 tx/s), and an upper bound, the postulated maximum value following the merge (3 ktx/s), is presented (cf. Appendix B).

⁷<https://docs.hedera.com/guides/mainnet/mainnet-nodes/node-requirements>

All estimates are based on the validator counts established earlier (cf. Section IV). The plot of the model function shown in Figure 1 visualizes the inverse relationship described earlier within the boundaries of the postulated throughput values (cf. Table III). It also provides a projection of energy consumption as a function of system load, based on the model presented earlier which predicts the number of validators as a function of system load. This projection is equally illustrated within the boundaries of the postulated throughput values, except in the case of the Tezos Blockchain for which no global model could be derived.

Based on this data, we can compare the energy consumption per transaction on two related systems: first the PoW cryptocurrency Bitcoin and second the VisaNet payment network (Figure 1). It becomes evident that the consumption of Bitcoin – overall and per-transaction – is at least three orders of magnitude higher than that of the highest consuming PoS system even under the most favorable assumptions. While the difference between PoS systems and VisaNet is less pronounced, it is evident that most of the former undercut the energy consumption of VisaNet in most configurations.

Platform	Global (kW)	Per transaction (kWh/tx)
Eth. 2.0[†]	1010.6 – 30 887.5	0.000 09 – 0.002 86
Eth. 2.0[‡]	1010.6 – 30 887.5	0.018 23 – 0.557 13
Algorand	6.2 – 189.3	0.000 17 – 0.005 34
Cardano	16.3 – 497.2	0.012 39 – 0.378 54
Polkadot	1.6 – 49.9	0.003 78 – 0.115 56
Tezos	2.2 – 67.1	0.000 36 – 0.010 96
Hedera	3.5 – 6.9	0.000 02 – 0.000 04
Bitcoin	3 373 287.7 – 34 817 351.6	360.393 00 – 3691.407 00
VisaNet	22 387.1	0.003 58

[†] High throughput projection

[‡] Low throughput projection

Table VI
GLOBAL POWER CONSUMPTION (I.E. THE NETWORK-WIDE CONSUMPTION OF THE DLT SYSTEMS UNDER CONSIDERATION AND VISANET) AND THE ENERGY CONSUMED PER TRANSACTION FOR CONTEMPORARY THROUGHPUT (SEE TABLE III)

Pronounced differences between PoS-based systems are equally evident from the results. We observe low energy demand per transaction in active permissioned DLT systems that are characterized by comparatively small numbers of validators and high throughput. Less active permissionless systems show a higher energy demand per transaction due to comparatively lower throughput and a high number of validators. This illustrates that not only for PoW [31] but also for PoS blockchains, “energy consumption per transaction” should not be the only metric considered for assessing the sustainability. Particularly when utility is not approximately proportional to throughput, total energy consumption may be a more appropriate key figure.

VI. DISCUSSION

A. Interpretations

These results can primarily be understood as a clear confirmation of the common opinion that the energy consumption of




Figure 1. The energy consumption per transaction is close to inversely correlated with throughput. For each system, the lower mark indicates the energy consumption under an optimistic validator hardware assumption while the upper mark indicates a pessimistic model. The consumption figures for Bitcoin and VisaNet are plotted for comparison (cf. Appendix D). For Ethereum 2.0, no throughput metrics are available.

PoW systems, especially Bitcoin, is excessive. Therefore, they can be interpreted as a strong argument for the modernization of PoW-based systems towards PoS. Ethereum is taking a commendable lead in this respect with the development of Ethereum 2.0. Furthermore, the results indicate that the energy consumption of different non-PoW blockchains is surprisingly divergent (e.g., by a factor of about 1×10^3 between the PoS system with the highest consumption and the one with the lowest). In absolute terms, however, the consumption rates of PoS-based systems are moderate and thus also much closer to the figures for traditional, centralized payment systems such as VisaNet.

The main reason why our model yields considerable divergence between PoS systems is the different number of validators. Specifically, in permissioned systems, energy consumption can be controlled through the ability to limit the number of validators on a network, so the permissioned network analyzed in this study is characterized by low energy consumption. However, this observation does not warrant conclusions such as that permissioned systems are necessarily less energy consumptive. Moreover, while in permissioned systems an operator can influence the number of nodes, it does not necessarily mean that that number must be lower.

Even if a reducing effect of permissioning on energy consumption could be stated with certainty, this should not be misinterpreted as an argument for increased centralization or an argument for permissioned networks over permissionless ones. This becomes obvious when considering a permissioned DLT system in extremis: such a system would consist of only a single validator node and would thus be effectively centralized. This hypothetical scenario shows that, if a permissioned paradigm is applied, close attention should be paid to system entry barriers

enforced through gatekeeping capabilities. If not, there is a risk of centralization, which may offer advantages in terms of energy consumption, but will negate the functional advantages of a decentralized paradigm. Of practical relevance is also the result that the selection of suitable validator hardware is central to energy consumption. Information regarding adequate hardware for validators is often inconsistent. Therefore, standardized recommendations should be put forward to help operators of validator nodes in selecting the most energy-efficient hardware configuration.

This study is only a first step towards quantifying the energy consumption of PoS systems. However, despite its limitations, it gives impetus to designers of decentralized systems by revealing the dependency between validator number, load, and hardware configuration. Our model can thus be used to determine the carbon footprint of a particular use case. It can furthermore prompt operators of validator nodes to carefully select suitable hardware.

B. Limitations

So far, we have used broad consumption ranges to model the energy consumption of individual validator nodes. While we are confident that the actual energy consumption is in fact within these ranges, underlying characteristics of different PoS protocols that might impact energy consumption, such as the accounting model, have been ignored. Second, while assuming that the electricity consumption of a validator node is independent of system throughput is well justified for the permissionless systems analyzed [27], permissioned systems that are designed to support high throughput may not warrant such assumption. While we have accounted for this by assuming more powerful hardware for permissionless hightps systems, more work is needed to understand permissioned blockchains'

energy consumption characteristics better. Moreover, the impact of different workloads on energy consumption should be considered; for example, simple payments transactions may have lower computational requirements when compared to other smart contract calls, but so far we have not distinguished between transaction types.

Further, while our model suggests that PoS systems can remain energy-efficient while scaling up to VisaNet throughput levels, there is no hard evidence in support of this argument, as no DLT-based system has experienced a sustained volume of this magnitude to date on the base level.

We ignored the possibility of achieving effectively higher throughput than the specified maximum through layer 2 (L2) solutions, such as the Lightning network or via optimistic and zero-knowledge (zk)-rollups that are receiving increasing attention.

Finally, although there are reasons to support its plausibility, the assumption that an affine function can be used to express the number of validators in terms of throughput is questionable. While we assume that it is applicable to Hedera, this might not be a justifiable assumption for other permissioned settings. The applicability of this model to other permissioned systems should therefore be more formally analyzed.

VII. CONCLUSION

The increasing popularity of DLT systems since the invention of Bitcoin, and with it the energy-intensive PoW consensus mechanism, has produced a variety of alternative mechanisms. PoS is a particularly popular alternative that is commonly assumed to be more energy efficient than PoW. In this paper, we tested this hypothesis using a mathematical consumption model that predicts expected energy consumption per transaction, as a function of network load. Applying this model to six different PoS-based DLT systems supports the hypothesis and suggests that their energy consumption per transaction is indeed at least two to three orders of magnitude lower than that of Bitcoin. Furthermore, we discover significant differences among the analyzed PoS-based systems themselves. Here, a permissioned system was found to consume significantly less energy per transaction than permissionless systems. This difference could be attributed to gatekeeping capabilities offered by permissioned systems.

These results can be understood as an urgent call for the modernization of PoW systems and a shift towards PoS, as well as a recommendation to practitioners to consider appropriate, energy-saving hardware. They are also intended to provide a basis for the future comparative study of the energy friendliness of PoS systems and to facilitate the development of more rigorous consumption models. Given the enormous challenges posed by climate change, avoiding unnecessary energy consumption needs to be a high priority. Our work shows that PoS-based systems can contribute to this and could even undercut the energy needs of traditional central payment systems, raising hopes that DLT can contribute positively to combatting climate change.

Future research should further develop and confirm these initial findings by improving the sophistication of the model and considering factors beyond network throughput, that may influence validator count. It should, furthermore, consider the network-wide energy consumption beyond validator nodes (i.e., by including all full nodes and auxiliary services) to arrive at a more holistic view of the overall energy consumption of DLT systems. Applying benchmarking frameworks [32] to measure the actual energy consumption might be particularly worthwhile in the context of permissioned systems that aim for high performance. In addition, analyzing the actual hardware configurations, instead of relying on rough estimates, might prove a worthwhile extension. Finally, future work should assess the effects of moving from a permissioned to a permissionless model.

ACKNOWLEDGEMENTS

We thank Michel Zade for comments that greatly improved the manuscript.

M.P. was supported by the University College London Centre for Blockchain Technologies. M.P. was also supported by Google Cloud via the Google Cloud Research Grant program. D.P. was supported by the Engineering and Physical Sciences Research Council [EP/L015234/1], the EPSRC Centre for Doctoral Training in Geometry and Number Theory (The London School of Geometry and Number Theory), University College London, and by Imperial College London.

AUTHOR CONTRIBUTIONS

Conceptualization: M.P., J.X., P.T., N.V. and J.I.I.; Data curation: M.P., J.S. and D.P.; Formal analysis: D.P.; Investigation: M.P., J.S., D.P., J.X. and J.I.I.; Methodology: M.P., J.S. and D.P.; Visualization: M.P. and D.P.; Writing – original draft: M.P.; Writing – review & editing: M.P., J.S., D.P., U.G., J.X., P.T., N.V. and J.I.I..

CONFLICT OF INTEREST

M.P. declares that he is bound by a confidentiality agreement that prevents him from disclosing his competing interests in this work.

ACRONYMS

CPU	central processing unit
DAG	directed acyclic graph
DLT	distributed ledger technology
L2	layer 2
LPoS	liquid proof-of-stake
NPoS	nominated proof-of-stake
PoS	proof-of-stake
PoW	proof-of-work
tps	transactions per second
UTXO	unspent transaction output
zk	zero-knowledge

REFERENCES

- [1] Y. Xiao, N. Zhang, W. Lou, and Y. T. Hou, "A survey of distributed consensus protocols for blockchain networks," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 2, pp. 1432–1465, 2020.
- [2] J. R. Douceur, "The Sybil attack," in *Proceedings of the 1st International Workshop on Peer-to-Peer Systems*, P. Druschel, F. Kaashoek, and A. Rowstron, Eds., ser. Lecture Notes in Computer Science, vol. 2429, Cambridge, MA, USA: Springer, 2002, pp. 251–260.
- [3] M. Platt, R. J. Bandara, A.-E. Drăgnoiu, and S. Krishnamoorthy, "Information privacy in decentralized applications," in *Trust Models for Next-Generation Blockchain Ecosystems*, ser. EAI/Springer Innovations in Communication and Computing, M. Rehman, D. Svetinovic, K. Salah, and E. Damiani, Eds., Forthcoming, Springer.
- [4] P. Tasca and C. J. Tessone, "A taxonomy of blockchain technologies: Principles of identification and classification," *Ledger*, vol. 4, Feb. 2019.
- [5] S. Nakamoto. (2008). "Bitcoin: A peer-to-peer electronic cash system," [Online]. Available: <https://bitcoin.org/bitcoin.pdf> (visited on 07/22/2021).
- [6] A. Back. (Mar. 1997). "A partial hash collision based postage scheme," [Online]. Available: <http://www.hashcash.org/papers/announce.txt> (visited on 07/27/2021).
- [7] J. Sedlmeir, H. U. Buhl, G. Fridgen, and R. Keller, "The energy consumption of blockchain technology: Beyond myth," *Business & Information Systems Engineering*, vol. 62, no. 6, pp. 599–608, Jun. 2020.
- [8] A. de Vries, "Bitcoin's growing energy problem," *Joule*, vol. 2, no. 5, pp. 801–805, May 2018.
- [9] L. Ismail and H. Materwala, "A review of blockchain architecture and consensus protocols: Use cases, challenges, and solutions," *Symmetry*, vol. 11, no. 10, p. 1198, Sep. 2019.
- [10] N. Lei, E. Masanet, and J. Koomey, "Best practices for analyzing the direct energy use of blockchain technology systems: Review and policy recommendations," *Energy Policy*, vol. 156, 2021.
- [11] A. Igumenov, E. Filatovas, and R. Paulavičius, "Experimental investigation of energy consumption for cryptocurrency mining," in *Proceedings of the 11th International Workshop on Data Analysis Methods for Software Systems*, J. Bernatavičienė, Ed., Druskininkai, Lithuania: Vilnius University Press, Nov. 2019, p. 31.
- [12] D. Saingre, T. Ledoux, and J.-M. Menaud, "BCTMark: A framework for benchmarking blockchain technologies," in *Proceedings of the 17th International Conference on Computer Systems and Applications*, Antalya, Turkey: IEEE, Nov. 2020, pp. 1–8.
- [13] C. A. Roma and M. A. Hasan, "Energy consumption analysis of XRP validator," in *Proceedings of the 2020 International Conference on Blockchain and Cryptocurrency*, IEEE, May 2020, pp. 1–3.
- [14] R. Cole and L. Cheng, "Modeling the energy consumption of blockchain consensus algorithms," in *Proceedings of the 2018 International Conference on Internet of Things and Green Computing and Communications and Cyber, Physical and Social Computing and Smart Data*, Halifax, NS, Canada: IEEE, Jul. 2018, pp. 1691–1696.
- [15] U. Gallersdörfer, L. Klaaßen, and C. Stoll, "Energy consumption of cryptocurrencies beyond Bitcoin," *Joule*, vol. 4, no. 9, pp. 1843–1846, Sep. 2020.
- [16] S. Küfeoglu and M. Özkuran, "Energy consumption of Bitcoin mining," University of Cambridge, Cambridge Working Paper in Economics 1948, 2019. doi: 10.17863/CAM.41230.
- [17] M. Zade, J. Myklebost, P. Tzscheuschler, and U. Wagner, "Is bitcoin the only problem? A scenario model for the power demand of blockchains," *Frontiers in Energy Research*, vol. 7, Mar. 2019.
- [18] J. Sedlmeir, H. U. Buhl, G. Fridgen, and R. Keller, "Ein Blick auf aktuelle Entwicklungen bei Blockchains und deren Auswirkungen auf den Energieverbrauch," German, *Informatik Spektrum*, vol. 43, no. 6, pp. 391–404, Nov. 2020.
- [19] H. Vranken, "Sustainability of Bitcoin and blockchains," *Current Opinion in Environmental Sustainability*, vol. 28, pp. 1–9, Oct. 2017.
- [20] G. Eshani, D. Rajdeep, R. Shubhankar, and D. Baisakhi, "An analysis of energy consumption of blockchain mining and techniques to overcome it," in *Proceedings of the International Conference on Computational Intelligence, Data Science and Cloud Computing*, V. E. Balas, A. E. Hassanien, S. Chakrabarti, and L. Mandal, Eds., ser. Lecture Notes on Data Engineering and Communications Technologies, vol. 62, Kolkata, India: Springer, 2021, pp. 783–792.
- [21] L. M. Powell, M. Hendon, A. Mangle, and H. Wimmer, "Awareness of blockchain usage, structure, & generation of platform's energy consumption: Working towards a greener blockchain," *Issues In Information Systems*, vol. 22, no. 1, pp. 114–123, 2021.
- [22] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich, "Algorand: Scaling byzantine agreements for cryptocurrencies," in *Proceedings of the 26th Symposium on Operating Systems Principles*, Shanghai, China: ACM, Oct. 2017, pp. 51–68.
- [23] C. Badertscher, P. Gaži, A. Kiayias, A. Russell, and V. Zikas, "Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability," in *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, 2018, pp. 913–930.
- [24] L. Goodman, *Tezos: A self-amending crypto-ledger*, 2014. [Online]. Available: https://cryptorating.eu/whitepapers/Tezos/position_paper.pdf (visited on 08/26/2021).

- [25] L. Baird. (Jun. 2016). “Swirls and sybil attacks,” [Online]. Available: <https://www.swirls.com/downloads/Swirls-and-Sybil-Attacks.pdf> (visited on 07/28/2021).
- [26] L. Hedera Hashgraph, *What is gossip about gossip?* [Online]. Available: <https://hedera.com/learning/what-is-gossip-about-gossip> (visited on 08/28/2021).
- [27] V. Buterin. (Apr. 2021). “Why sharding is great: Demystifying the technical properties,” [Online]. Available: <https://vitalik.ca/general/2021/04/07/sharding.html> (visited on 07/28/2021).
- [28] M. Platt, *Cardano throughput and stake pool sizes*, Dataset, 2021. DOI: 10.17632/4JV2WMWRC5.1.
- [29] A. Jaiantilal, Y. Jiang, and S. Mishra, “Modeling CPU energy consumption for energy efficient scheduling,” in *Proceedings of the 1st Workshop on Green Computing*, Bangalore, India: ACM, 2010, pp. 10–15.
- [30] Z. Zhou, J. H. Abawajy, and F. Li, “Analysis of energy consumption model in cloud computing environments,” in *Advances on Computational Intelligence in Energy*, Springer, 2019, pp. 195–215.
- [31] N. Carter, *How much energy does bitcoin actually consume?* Harvard Business Review, 2021. [Online]. Available: <https://hbr.org/2021/05/how-much-energy-does-bitcoin-actually-consume> (visited on 08/26/2021).
- [32] J. Sedlmeir, P. Ross, A. Luckow, J. Lockl, D. Miehle, and G. Fridgen, “The DLPS: A framework for benchmarking blockchains,” in *Proceedings of the 54th Hawaii International Conference on System Sciences*, 2021, pp. 6855–6864.

APPENDIX

A. Validator Metrics

Chain	Source	Metric	Obs. Period	Value
Ethereum 2.0	https://beaconcha.in/charts	Number of active validators	5/7/2021	183 753
Algorand	https://metrics.algorand.org/	Number of nodes	12/8/2021	1126
Cardano	https://cardanoscan.io/	Number of stake pools	11/8/2021	2958
Polkadot	https://polkadot.subscan.io/validator	Number of validators	5/7/2021	297
Tezos	https://tzstats.com/bakers	Number of bakers	12/8/2021	399
Hedera	https://docs.hedera.com/guides/mainnet/mainnet-nodes	Numbers of mainnet nodes	13/8/2021	21s

Table VII

SOURCES FOR DATA ON CONTEMPORARY VALIDATOR MACHINE COUNT

Chain	Source	Metric	Obs. Period	Value
Polkadot	https://web.archive.org/web/*/https://stakers.info/	Number of validators	27/2/2021	297
Tezos	https://api.tzstats.com/explorer/cycle/324	Number of bakers	5/2/2021	430
Algorand	https://metrics.algorand.org/	Number of validators	5/7/2021	1298
Hedera	https://docs.hedera.com/guides/mainnet/mainnet-nodes	Number of validators	5/7/2021	20
Hedera	https://github.com/hedera-docs/commits/master/mainnet/mainnet-nodes/README.md	Number of validators	7/7/2020–26/8/2021	20

Table VIII

SOURCES FOR DATA ON HISTORIC VALIDATOR MACHINE COUNT

B. Throughput Metrics

Chain	Source	Metric	Obs. Period	Value
Algorand	https://algorandexplorer.io/	Average transaction volume	16/7/2021-12/8/2021	9.845 tx/s
Cardano	https://explorer.cardano.org/en	Number of transactions in epoch	Epoch 282 (3/8/2021-8/8/2021)	157 622 tx
Polkadot	https://polkadot.subscan.io/extrinsic	Mean of the lowest and the highest daily transaction volume	5/6/2021-5/7/2021	0.1200 tx/s
Tezos	https://tzstats.com/	Average number of transactions per second	13/7/2021-12/8/2021	1.700 tx/s
Hedera	https://hedera.com/dashboard	Transaction volume by network service	13/8/2021	48.20 tx/s

Table IX

SOURCES FOR DATA ON CONTEMPORARY THROUGHPUT

Chain	Source	Metric	Obs. Period	Value
Algorand	https://algorandexplorer.io/	Transactions per second	2/6/2021-2/7/2021	11.5 tx/s
Tezos	https://messari.io/asset/tezos	Average number of transactions per second	6/1/2021-5/2/2021	0.4 tx/s
Hedera	https://hedera.com/dashboard	Transactions per second	5/7/2021	44.6 tx/s
Hedera	https://app.dragonglass.me/hedera/home	Transactions per second	8/2020-8/2021	-

Table X
SOURCES FOR DATA ON HISTORIC THROUGHPUT

Chain	Source	Metric	Value
Ethereum 2.0	https://twitter.com/VitalikButerin/status/1277961594958471168	Transactions per second with Ethereum 1 as data layer	3000 tx/s
Algorand	https://www.algorand.com/resources/blog/algorand-2021-performance	Current maximum transactions per second	1000 tx/s
Cardano	https://vacuumlabs.com/blog/life-vacuum/what-we-love-about-cardano-a-technical-analysis	Maximum theoretical throughput	257 tx/s
Polkadot	https://twitter.com/gavofyork/status/1255859146127179782	Sustained transactions per second	1000 tx/s
Tezos	https://blockfyre.com/tezos-xtz/	Transactions per second	40 tx/s
Hedera	https://hedera.com/hbar	Transactions per second	10 000 tx/s

Table XI
SOURCES FOR DATA ON MAXIMUM THROUGHPUT

Bound	Source	Metric	Obs. Period	Value
Lower	https://etherscan.io/	Throughput of Ethereum 1	24/7/2021	15.40 tx/s
Upper	https://twitter.com/VitalikButerin/status/1277961594958471168	Postulated maximum transactions per second	-	3000 tx/s

Table XII
SOURCES FOR THROUGHPUT ESTIMATES FOR ETHEREUM 2.0

C. Validator Energy Consumption

Hardware	Source	Metric	Value
Raspberry Pi 4	https://www.tomshardware.com/uk/reviews/raspberry-pi-4	Power consumption when idle	3.4 W
Raspberry Pi 4	https://www.tomshardware.com/uk/reviews/raspberry-pi-4	Power consumption under load	7.6 W
Dell PowerEdge R730	https://i.dell.com/sites/csdocuments/CorporateComm_Docs/en/carbon-footprint-poweredge-r730.pdf	Typical yearly energy consumption	1473.5 kWh
Hewlett Packard Enterprise ProLiant ML350 Gen10	https://www.specc.org/power_ssj2008/results/res2019q2/power_ssj2008-20190312-00899.html	Power consumption under 80% load	328 W

Table XIII
SOURCES FOR DATA ON HARDWARE ENERGY CONSUMPTION

D. Comparison Values

System	Source	Metric	Obs. Period	Value
Bitcoin	https://cbeci.org/	Theoretical lower bound of annualized power consumption	11/8/2021	29.55 TW h
Bitcoin	https://cbeci.org/	Theoretical upper bound of annualized power consumption	11/8/2021	305 TW h
Bitcoin	https://www.blockchain.com/charts/transactions-per-second	Transactions per second	30 day average on 11/8/2021	2.620 tx/s
VisaNet	https://usa.visa.com/content/dam/VCOM/global/about-visa/documents/visa-2020-esg-report.pdf	Approximate total energy consumption of the Visa corporation	2020	706 000 GJ
VisaNet	https://usa.visa.com/run-your-business/small-business-ess-tools/retail.html	Transactions per day	8/2010	150 Mtx/d

Table XIV
SOURCES FOR DATA ON REFERENCE SYSTEMS

About UCL CBT

The UCL CBT is the first centre globally to actively focus on blockchain-related research on the adoption and integration of Blockchain and Distributed Ledger Technologies into our socio-economic system.

The unique characteristics of the CBT at UCL provides a cross-sectoral platform connecting expertise and drawing knowledge from eight UCL departments centrally in one place. The CBT is a centre of excellence fostering open dialogue between industry players and sharing expertise and resources. It is a neutral think tank providing consultancy services to industry members, dedicated knowledge-transfer activities and cutting-edge in-house solutions.

For engagement outside of the academic world, the CBT's activities have been tailored to industry and policymakers' needs. The UCL CBT draws on its world-leading academic expertise to produce blockchain solutions for industry, start-ups and regulators. With a community of over 247 Research & Industry Associates and Industry Partners, it is the largest Academic Blockchain Centre in the world.

Notable Work

- The CBT released a report on the current adoption of DLT in global physical supply chains. The report featured an analysis of over 100 different projects taking place all over the world in the Grocery, Pharmaceutical and Fashion industries. Access the report [here](#).
- The CBT is a founding member of the [Covid Task Force](#) alongside The International Association for Trusted Blockchain Applications (INATBA) and the European Commission. The task force is convening key players in the global blockchain ecosystem to identify deployable technology solutions that address governmental, social, and commercial challenges caused by COVID. As well as identifying solutions, the Task Force will work to expedite their deployment.
- The CBT successfully funded nine research proposals that investigated topics including stable coin policy, smart contract innovation, blockchain economics and blockchain governance models. Research teams who were funded were made up of individuals from a variety of academic and industry organisations. Learn more about the projects [here](#).
- The CBT launched the Block-Sprint hackathon to promote DLT innovation in the financial services sector. Over 100 individuals took part in both the 2019 and the 2020 edition forming teams made up of industry practitioners, academics, and students. Learn about the winners and innovate ideas that were generated in the hackathon [here](#).

About the Discussion Paper Series

The *UCL CBT Discussion Paper* is published on a quarterly basis featuring the latest developments in the blockchain and DLT space. The aim of the discussion paper series is to share recent developments and state-of-the-art solutions on blockchain and DLT of researchers from an interdisciplinary background with the CBT community. All accepted submissions are available in the CBT paper database.

The submissions are circulated among the members of the UCL CBT Editorial Board, led by the Scientific Director so that the results of the research receive prompt and thorough professional scrutiny.

If you are interested in submitting a paper to be included in forthcoming editions, please visit our website [here](#) to see what the latest theme and criteria for submission are.

UCL Centre for Blockchain Technologies

<http://blockchain.cs.ucl.ac.uk/>

UCL Computer Science
Malet Place
London WC1E 6BT
United Kingdom

