DEcon: Decentralised Economics for Distributed Ledgers

Rowdy Chotkan
R.M.Chotkan-1@tudelft.nl
Delft University of Technology
Delft, Netherlands

Johan Pouwelse
J.A.pouwelse@tudelft.nl
Delft University of Technology
Delft, Netherlands

Abstract

Existing distributed ledgers are suffering from emergent centralization. A key driving force behind this centralization is the popularity of mining and staking pools. Nodes that are part of such a pool are virtually the same entity and, as such, there effectively exist few entities within even major blockchain networks. Mining and staking pools incentivize nodes by consistent payouts, lower barriers of entry, and most importantly acknowledgement of partial work. This last benefit is crucial as it enables consistent payouts and fosters efficiency due to divide-and-conquer in Proof-of-Work blockchains. As such, their existence is favourable for the health of the network and, hence, should not be disincentivized. In order to address concerns regarding centralization whilst still harbouring the benefits of the pooling of resources, this work proposes *DEcon*: a decentralized economic model for distributed ledger technologies. DEcon rewards nodes in distributed ledgers based on all useful work they perform, both on layer 0 and layer 1, allowing for fair and consistent payouts. This is performed through local acknowledgements of useful work that nodes perform. These acknowledgements also occur for invalid tested hashes. Allowing clients to be compensated for their partial work in finding blocks. In the process, DEcon can be considered a fully decentralised mining pool. Nodes are rewarded based on a competitiveness function.

CCS Concepts: • **Do Not Use This Code → Generate the Correct Terms for Your Paper**; *Generate the Correct Terms for Your Paper*; Generate the Correct Terms for Your Paper; Generate the Correct Terms for Your Paper.

Keywords: Do, Not, Us, This, Code, Put, the, Correct, Terms, for, Your, Paper

DICG '23, June 03–05, 2018, Woodstock, NY
© 2023 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/XXXXXXX.XXXXXXX

ACM Reference Format:
Rowdy Chotkan and Johan Pouwelse. 2023. *DEcon:* Decentralised Economics for Distributed Ledgers. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation emai (DICG '23).* ACM, New York, NY, USA, 7 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

We present the first economic model for distributed ledger technology (DLT) with *full decentralisation*. Centralization, of course, is detrimental to the overall health of the decentralized system. More specifically, centralization in systems can lead to security and privacy issues, power imbalances, and political and economic issues [4]. The effects of which, we have observed in practice with e.g., MEV [9]. As such, due to the emergent and rampant centralisation of power within layer 1 of prominent blockchain technologies, we deem it a necessity to devise a novel approach to sustain decentralization.

The key question is what the driving force behind that centralisation—or rather *re*centralization—is. We believe that mining and staking pools have created an ecosystem facilitating this centralisation. In these pools, clients pool their power, be it hash rate or stake. These pools are then coordinated through a central orchestrator, allowing them to achieve an optimal mining strategy by applying a divide-and-conquer method by splitting up the search space or by raising the odds through the pooled stake. Especially for smaller clients, who prefer consistent payout, joining a mining pool is the most rational approach, as otherwise, their payout may take years. Prominent ledgers based on Proof-of-Stake are even more exclusive: staking requires a minimum stake, excluding stakes with limited funds [7]. For reference, the minimum stake as of writing for Ethereum converts to roughly 60 thousand USD.

The degree of decentralization of a blockchain can be represented by the distribution of mining and staking power, an indicator which is quantifiable by the Nakomoto coefficient [13]. Mining pools—and staking pools equally—have severely limited the number of major entities within blockchains [14]. Due to pooling, all nodes are effectively part of the same entity, due to them exhibiting equivalent behaviour orchestrated by a central coordinator, leading to a decrease in the degree of decentralization in a system. For instance in Bitcoin, over 50% of the hash rate is controlled by two entities [3]. In Ethereum, the problem is currently less prominent (e.g., the largest staking pool possesses roughly 14% of the network [1]), however, concerns over rising centralization imposed by staking pools have been raised extensively (e.g.,[5, 6, 8]).

Decentralisation has been proven to be extremely difficult. Only Bitcoin and BitTorrent are examples of protocols with enduring decentralisation. The email protocol started as a decentralised protocol, but spam and security in general have eroded decentralisation. Running your own personal email server is no longer viable [REF?]. Similar with IPFS, this decentralised protocol only has a single durable gateway operated by Cloudfare. BitTorrent also has seen such erosion of decentralisation with big-gets-bigger websites for content discovery. Bitcoin decentralisation is also eroding with dominant mining pools. Therefore, no example exists of a durable decentralised protocol which can serve as a blueprint for decentralised economics.

We refer to the pooling of resources as *indirect participation*: nodes do not verbatim act as full nodes, rather they serve the mining or staking pool with respect to the transaction and block sharing and partial mining. We believe that the need for mining pools goes beyond the consistent payouts, rather the intrinsic value they provide is the rewarding of partial work carried out: for instance, in Proof-of-Work (PoW) blockchains, invalid hashes for the next block are rewarded. As such, we raise ourselves the question *how to reward direct participation as well as indirect participation?*. By answering this question, we aim to devise a new economic paradigm in distributed ledger technologies that goes beyond the consensus layer. In turn, nodes are rewarded for *any* work—dubbed utility— they create or generate for the network. Such a universal mechanism is capable of overcoming the centralisation risks associated with mining and staking pools by essentially creating a single, globally shared, mining pool. Clients are thus not only rewarded for their efforts in reaching consensus but also on layer 0: they are rewarded based on their utility generated through networking (e.g. sharing mempools, bootstrapping nodes, and availability). We dub this reward mechanism, the **competitiveness function**.

We present a novel approach to Web3 economics, based upon academically-pure decentralisation. The *conservation of centralisation* paradox is central to our work. It is defined as: removing centralisation is an illusion, it will re-appear. Our unique approach is to take long-enduring decentralisation as the core objective. We are the first to explicitly identify the lack of any past success around academically pure decentralisation in Web3. We present the first meticulously designed Web3 architecture which avoids re-centralisation.

The reasoning that this new economic system does not try to absolve the pooling of resources, lies in the additional value that mining and staking pools generate: they enable a higher throughput of transactions due to the sharing of mempools and, most importantly, they are able to enforce a divide-and-conquer mechanism therefore increasing the efficiency of PoW blockchains. Our new mechanism, when globally applied, will therefore be able to make any PoW blockchain more efficient and in turn lower their overall performance needs. Furthermore, nodes in any PoS blockchain will be compensated not only by chance but also by their essential networking activities.

2 Introduction Old

We present the first economic model for distributed ledger technology (DLT) with *full decentralisation*. We see a need for a new economic model due to emergent centralisation in DLT. Web2 economics is defined by *centralisation*. Web2 companies aim to get-big-fast through acquisitions. For instance, the message company Whatsapp famously employed 32 engineers when it was bought by Facebook for $16 billion [1]. Over the last 10 years the 5 largest tech firms have made over 400 acquisitions globally, according to a UK government report[2]. This UK report also recommends to "enforce a clear set of rules to limit anti-competitive actions by the most significant digital platforms". "Acquisitions of big tech firms are systematically not investigated" by regulators[3] Acquisitions lead to further centralisation. We currently observe centralisation of entire markets into a single company (search, social, e-commerce, etc.). The *killer acquisition* theory states that Big Tech acquires startups with the potential to become a competitive threat. The US National Bureau of Economic Research states Big Tech "tends to acquire younger targets is consistent with the concern of "killer acquisitions""[4]. Bill Gates from Microsoft wrote a damaging email in August 1997 – "Do we have a clear plan on what we want Apple to do to undermine SUN?" [5]. Decades later, avoiding anti-competitive emails seems to be part of "Legal 101" training for employees. Avoiding certain words is "the one big thing I remember from all that Legal training. :-)"[6] according to an internal Big Tech email published in September 2023 by the US Department of Justice in their anti-competition trail.

Decentralisation has been proven to be extremely difficult. Only Bitcoin and BitTorrent are examples of protocols with enduring decentralisation. The email protocol started as a decentralised protocol, but spam and security in general

[1] ttps://www.wsj.com/articles/facebook-to-buy-whatsapp-for-16-billion-1392847766

[2] https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/785547/unlocking_digital_competition_furman_review_web.pdf

[3] https://www.wiwi.uni-bonn.de/bgsepapers/boncrc/CRCTR224_2020_147v2.pdf

[4] https://www.nber.org/system/files/working_papers/w29642/w29642.pdf

[5] https://www.justice.gov/atr/file/704856/download

[6] https://www.justice.gov/d9/2023-09/416652.pdf

have eroded decentralisation. Running your own personal email server is no longer viable [REF?]. Similar with IPFS, this decentralised protocol only has a single durable gateway operated by Cloudfare. BitTorrent also has seen such erosion of decentralisation with big-gets-bigger websites for content discovery. Bitcoin decentralisation is also eroding with dominant mining pools. Therefore, no example exists of a durable decentralised protocol which can serve as a blueprint for Web3.

Centralisation has a long history. Every industry is concentrated in fewer and fewer hands. In a broader historical economical context, the centralisation of firms has been ongoing for centuries. For instance, see the shipbuilding centralisation in the Venetian Arsenal from 1172[7]. Digitisation further accelerates centralisation and gives rise to large monopolistic firms[8]. Monopoly firms erode the principle of market competition. Instead of competition *within* markets, they compete *for* markets. For instance, Facebook (Meta) used their existing user base to attack the de-facto X (Twitter) monopoly on microblogging owned by Elon Musk[9].

Decentralisation is seen by many as a cure for the emergence of natural Big Tech monopolies. The cardinal principle is forcing collaborative relations between competing firms. Firms have a strong incentive to lock users into their ecosystem. Regulators should enforce that their citizens can easily switch between products. Consumer well-being dictates that products need to work together, support open standards, be interoperable, and compete directly on an equal playing field. However, we argue that government action is consistently too slow in dynamic digital markets and see adversarial interoperability as the defining feature for Web3[10]. Web3 should be compatible with Web2, until a court re-defines such nonprofit interoperability as an *illegal innovation*. We believe this is what *real* Web competition looks like. In the absence of competitive offerings from rival firms, citizens cooperate to craft an alternative without asking permission.

Web3 is a leaderless decentralised movement with the dream of providing alternatives to Big Tech. Yet, no viable economic model of Web3 has been presented to date. Only speculative tokens-based approaches have been proposed. Such tokens, nano-currencies [12] or proof-of-x functions [11] lack the inherent utility that, for instance, gold possesses.

We present a novel approach to Web3 economics, based upon academically-pure decentralisation. The *conservation of centralisation* paradox is central to our work. It is defined as: removing centralisation is an illusion, it will reappear. ADD See centralisation in various markets: REF https://cacm.acm.org/magazines/2021/3/250710-the-decline-of-computers-as-a-general-purpose-technology/fulltext

Our unique approach is to take long-enduring decentralisation as the core objective. We are the first to explicitly identify the lack of any past success around academically-pure decentralisation in Web3. We present the first meticulously designed Web3 architecture which avoids re-centralisation.

3 Problem Description - draft

We define the central problem of DLT economics as: *how to organise trustworthy economic activities in trustless decentralised networks in which strangers interact directly, removing the need for any intermediary, abolishing all central authorities while still providing protection against misreporting, slandering, fake identities, Sybil attacks, and fraud in general?*

Our new economic mechanism must provide flexibility in terms of utility generated by nodes within the network, whilst simultaneously not lessening the performance of the blockchain. In other words, the protocol must make any amount of work profitable whilst still incentivizing larger clients to perform the bulk of the work. We construct the following set of requirements:

- No central orchestrators: there exist no single client within the economy that possesses more power on a protocol level.
- Unstoppable: the protocol cannot be stopped by anyone. It runs for as long as the underlying blockchain exists.
- Sybil resistant: Sybils must not posses a greater advantage than the benefit they provide.

We raise ourselves the challenge of devising an economy purely based on the notion of *competitive fitness*: the more utility a node provides, the more reward they shall reap. The utility u_i of a node i can be defined by the sum of its performed work U_i:

$$u_i = \sum_{x \in U_i} x$$

Using said scoring, nodes must be rewarded proportionally to the amount of utility provided by the network. This must be performed in such a manner that:

- Large nodes must still receive the bulk of the reward.
- Small nodes must receive at least as much reward as when pooling resources.
- Fees must not be heavily impacted by the fitness function.

The fitness function itself can be defined as follows. Whenever a new block is formed, the fitness function is run over the set of existing nodes:

$$f(X, n) = \{c..., c..., \cdots, c...\}$$

[7] https://www.academia.edu/37935615/BONDIOLI_Mauro_The_Arsenal_of_Venice_and_the_Art_of_Building_Ships

[8] https://mudancatecnologicaedinamicacapitalista.files.wordpress.com/2019/02/platform-capitalism.pdf

[9] https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4500582

[10] https://www.eff.org/deeplinks/2019/06/adversarial-interoperability-reviving-elegant-weapon-more-civilized-age-slay

Where n is a system-dependent variable determining the size of the reward pool. I.e., n determines the amount of nodes that receive a partial reward.

4 Old Problem Description

Our aim is to devise a mechanism that is able to unambiguously and transparently donate to nodes within a Web3 ecosystem and therefore in any type of decentralised network. We dub this a *Web3 beneficiary function*. When run, this function should select a set amount of contenders from the network and reward them in the form of cryptocurrency, the amount of which is set by the donor. The chance of receiving this donation is relative to the amount of useful work—or *utility*— a client has put in. The end goal of this function is an ecosystem that rewards those based on useful work that has been performed for the network as opposed to those able to optimise their mining strategy. This useful work thus spans further than the mining of blocks. Realising this function, however, is not a trivial problem. We identify the following issues.

Client selection — First, the most prominent issue lies before the distribution of rewards. It is non-trivial to select those clients that are deemed useful to the network. With blockchains having no central orchestrators, registration of useful work is not a straightforward feat. One might be tempted to simply reward clients based on the number of blocks they mined, though this would still simply reward those with the most mining power. As such, this system would not net a more fair situation than currently exists. The essence of this issue, thus, lies in the quantification of useful work.

Network views — The aforementioned issue mostly comes into play when donating clients do not have a proper network view. When clients have a too-small network view to properly assess nodes within the network, they will revert to the knowledge of others. We refer to this phenomenon in its extreme form as the *bootstrapping problem*. In the bootstrapping problem, new clients have no network view at all and must thus entirely rely on existing nodes. In a practical setting, this would d be transacting clients within a blockchain network that relies on full nodes to handle their transactions. In the instance that no full network view exists, the responding clients might give biased results for their personal gain. **Sybils** — A selection of clients poses an additional challenge. A malicious actor can consistently rank its own Sybils higher than actual useful nodes within the network. If not addressed, unhelpful clients could receive donations whilst not actively maintaining the network.

Transactions — Second, the distribution of the donation or reward itself poses a challenge. When distributing rewards, a trade-off has to be made between the number of clients to be rewarded and the corresponding transaction fees. Of course, a transaction's size increases with the number of recipients and, therefore, transaction costs. If the number of recipients is 1, the function becomes practically equivalent to the Proof-of-Stake algorithm as it will become a lottery—although weighted—between those who have performed some work for the network. Additionally, depending on the quantity of transactions, it could take a tremendous time for some nodes to be rewarded for their efforts.

4.1 System Model

We consider a permissionless network of N clients which communicate in a peer-to-peer fashion. We assume that communication channels are established in a decentralised manner and that each client has a cryptographic key pair which allows for the authentication and verification of messages. For simplicity's sake, one can assume that clients communicate as nodes within the Bitcoin network.

The network itself can be represented by the graph $G = (V, E)$, where $V = \{v_0, v_1, \cdots, v_n\}$ represents the nodes within the network and $E = V$ rep

We want to devise a donation mechanism that rewards those that have contributed useful work for the network because that is a direct indicator for their level of maintenance towards the network.

5 System Model

We consider a network of nodes. This network is decentralized, egalitarian, and permissionless. We aim to create a mechanism that allows for unbiased donations to nodes within this network, based on the effective work they have performed, dubbed *utility*. This utility can be generated through any type of action that is performed for the benefit of the network. E.g., gossiping blocks and transactions in a blockchain or seeding data in P2P file-sharing.

The aim of this beneficiary function is to provide a means for regular users to donate to a decentralized network or project.

We present our set of properties which we deem necessary in order to achieve the set-out alternative incentive mechanism. This set of principles will enable us to devise a mechanism that is both fair to nodes within the network and profitable, whilst continuing to incentivise nodes to improve the services they provide.

Transparency — The act of donation must be a transparent process. It must be clear to the network how the set of receiving clients was formed. This is essential in creating a verifiable donation system in which we allow for the aforementioned proof-of-donations.

Variability — Verifiability enables the network to ensure that a client has donated to the network, enabling them to provide the clients with e.g. additional privileges.

graphics/tx_size.pdf

Figure 1. Transaction size

Incetivization — A crucial component of the mechanism is, of course, incentivization. It creates an ecosystem in which miners continue to progress the blockchain and provide services for clients. However, this incentivization, considering the new interplay of donation-based incentives, must go both ways.

Anonimity — Clients must have the ability to remain pseudo-anonymous in their donations.

Proportional compensation — Nodes must be proportionally compensated relative to the amount of actual work they have put into the network.

Minimal overhead — Our protocol must not impose heavy overhead onto the network, as doing so would defeat the purpose.

6 Solution Overview

We define our function as follows.

7 Results

In order to validate the effectiveness of our proposed beneficiary mechanism, we ran three different experiments.

We base our results on a Bitcoin experiment, simulating 10319 Bitcoin nodes—the number of reachable Bitcoin nodes reported by [2]. Using their peer index ranking mechanism, which scores peers numerically from 0 to 10 based on its properties and network metrics, we crafted a dataset representing nodes ranked by MeritRank [10]. Using this dataset, we ran an experiment in which fees are distributed amongst nodes, in which each ranking is translated to the probability of receiving said fees.

?? showcases the simulated network after 10^5, 10^6, 10^7, and 10^8 transactions, respectively. In these graphs, the percentage of work performed by clients with respect to the entire network is visualized together with the relative percentage of donations said clients have received. These results showcase that even with a relatively small number of transactions, our mechanism can distribute fees across participating nodes in a fair manner. To put these numbers into perspective: Bitcoin has had days with 700 thousand transactions[11].

Figure 4 showcases the impact that reliance on committees has on the convergence on the distribution of rewards. Here, committee members are modelled as individual nodes with a differing network view. I.e., nodes are assigned a random uniformly distributed factor score for each other node from $[-1, 1]$. As a consequence, no two nodes have the exact same score for other nodes, though well-performing nodes are consistently ranked higher and ill-performing nodes consistently lower. In this simulation, we report on the distribution of fees after 10^5 transactions. As visible, our mechanism is able to achieve a fair distribution with relatively small committee amounts. E

8 Related Work

9 Future work

10 Conclusion

Acknowledgments

This work was funded by NWO/TKI grant BLOCK.2019.004.

References

[1] beaconcha.in. 2023. Pool Distribution. https://beaconcha.in/pools

[2] Bitnodes. 2023. Bitcoin Network Leaderboard. https://bitnodes.io/nodes/leaderboard/

[3] Blockchain.com. 2023. Hashrate Distribution. https://www.blockchain.com/explorer/charts/pools

[4] Balázs Bodó, Jaya Klara Brekke, and Jaap-Henk Hoepman. 2021. Decentralisation: A multidisciplinary perspective. *Internet Policy Review* 10, 2 (2021), 1–21.

[5] Carlo Campajola, Raffaele Cristodaro, Francesco Maria De Collibus, Tao Yan, Nicolo' Vallarano, and Claudio J Tessone. 2022. The evolution of centralisation on cryptocurrency platforms. *arXiv preprint arXiv:2206.05081* (2022).

[6] Dap. 2022. Ethereum's centralization dilemma through Lido staking. https://www.suresats.com/post/ethereum-s-centralization-dilemma-through-lido-staking

[7] Hans Gersbach, Akaki Mamageishvili, and Manvir Schneider. 2022. Staking pools on blockchains. *arXiv preprint arXiv:2203.05838* (2022).

[8] Jack Inabinet. 2023. Leave lido alone. https://www.bankless.com/lido-did-nothing-wrong

[9] MEV Watch. 2023. MEV Watch. https://www.mevwatch.info

[10] Bulat Nasrulin, Georgy Ishmaev, and Johan Pouwelse. 2022. Meritrank: Sybil tolerant reputation for merit-based tokenomics. In *2022 4th Conference on Blockchain Research & Applications for Innovative Networks and Services (BRAINS)*. IEEE, 95–102.

[11] Giang-Truong Nguyen and Kyungbaek Kim. 2018. A survey about consensus algorithms used in blockchain. *Journal of Information processing systems* 14, 1 (2018).

[12] Wesley Sheh and Joel Nishimura. 2023. Why You Should Be Able to Make Your Own Individualized, Digital Nano-Currency. *Commun. ACM* 66, 8 (2023), 35–38.

[13] Balaji S. Srinivasan and Leland Lee. 2017. Quantifying Decentralization. https://news.earn.com/quantifying-decentralization-e39db233c28e

[11] https://ycharts.com/indicators/bitcoin_transactions_per_day

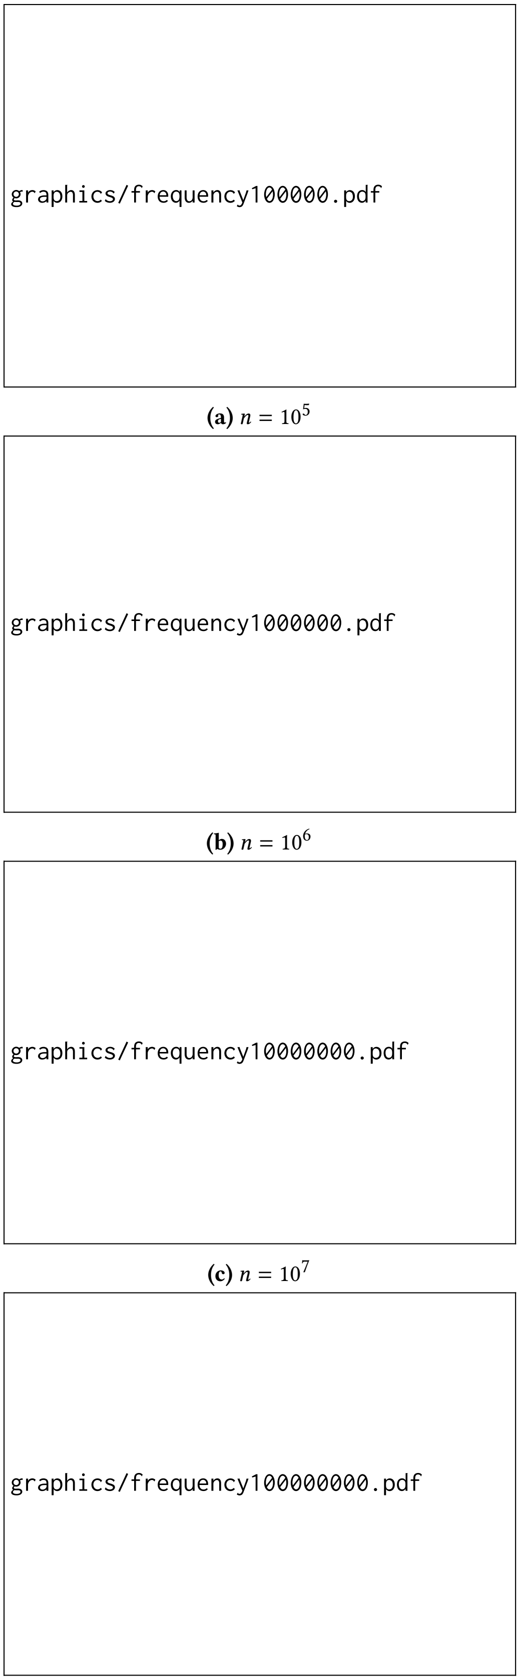
graphics/frequency100000.pdf
graphics/frequency1000000.pdf
graphics/frequency10000000.pdf
graphics/frequency100000000.pdf

(a) $n = 10^5$

(b) $n = 10^6$

(c) $n = 10^7$

(d) $n = 10^8$

Figure 2. Donations received

Figure 3. Converge of fee distribution function

[14] Zibin Zheng, Shaoan Xie, Hong-Ning Dai, Xiangping Chen, and Huaimin Wang. 2018. Blockchain challenges and opportunities: A survey. *International journal of web and grid services* 14, 4 (2018), 352–375.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009

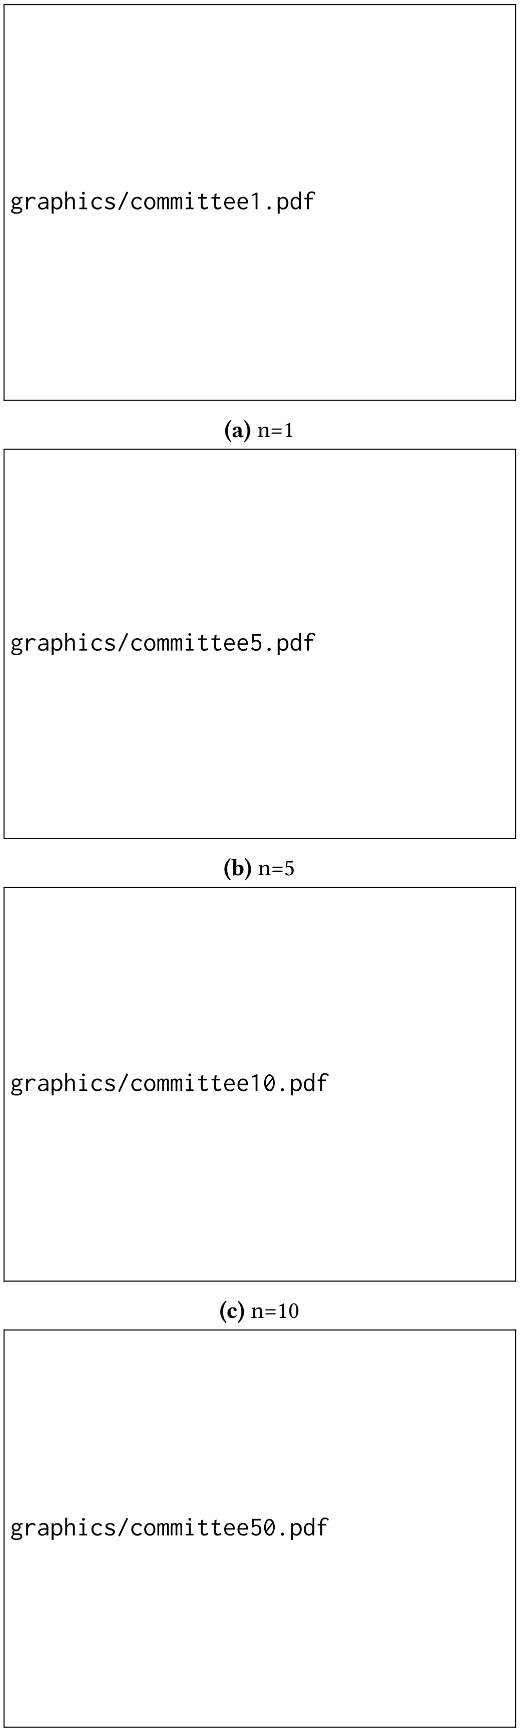

(a) n=1

(b) n=5

(c) n=10

(d) n=50

Figure 4. Donations received using committees