Decentralizing Computation with Edge Computing: Potential and Challenges

Aleksandr Zavodovski
Lorenzo Corneo
Andreas Johnsson
Uppsala University
firstname.lastname@it.uu.se

Nitinder Mohan
Technical University of Munich
mohan@in.tum.de

Suzan Bayhan
University of Twente
s.bayhan@utwente.nl

Pengyuan Zhou
University of Science and Technology of China
pyzhou@ustc.edu.cn

Walter Wong
Jussi Kangasharju
University of Helsinki
firstname.lastname@helsinki.fi

ABSTRACT

Edge computing promises to bring computation close to the end-users to support emergent applications such as virtual reality. However, the computational capacity at the edge of the network is currently limited. To become a pervasive paradigm, edge computing needs highly dispersed decentralized deployments, that, contrary to cloud, cannot benefit from economies of scale. In this situation, crowdsourcing appears attractive – there are plenty of computing devices at the disposal of the general public, and these devices are located exactly where computing power is needed the most – at the edge of the network. Crowdsourcing has been a success maker for scientific computing projects, e.g., SETI@home, or distributed ledger systems empowering decentralized finance. However, as of now, there is no crowdsourced system that addresses the needs of edge computing. In this position paper, we aim to identify the causes of this shortcoming, analyze the potential ways to overcome it, and outline future directions.

CCS CONCEPTS

• **Networks** → **Cloud computing**; • **Information systems** → **Crowdsourcing**.

KEYWORDS

Edge computing, cloud computing, crowdsourcing, blockchain

ACM Reference Format:
Aleksandr Zavodovski, Lorenzo Corneo, Andreas Johnsson, Nitinder Mohan, Suzan Bayhan, Pengyuan Zhou, Walter Wong, and Jussi Kangasharju. 2021. Decentralizing Computation with Edge Computing: Potential and Challenges. In *Interdisciplinary Workshop on (de)Centralization in the Internet (IWCI '21), December 7, 2021, Virtual Event, Germany*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3488663.3493689

IWCI '21, December 7, 2021, Virtual Event, Germany

ACM ISBN 978-1-4503-9138-2/21/12…$15.00
https://doi.org/10.1145/3488663.3493689

1 BACKGROUND & MOTIVATION

Edge computing, initially introduced via cloudlets [27], was primarily aimed to reduce the latency that end-users of computationally-limited devices experienced while accessing the cloud. Later, with the proliferation of Internet-of-Things (IoT), fog computing [6] was introduced to enable on-site aggregation, preprocessing, and data analysis for IoT flows. Currently, edge computing is a commonly-used term to denote an enabling technology for applications with stringent latency requirements, as well as for IoT. In the latter case, the motivation is to save the network and cloud from being overwhelmed by raw data flows and preserving user privacy.

As cloudlets were introduced already in 2009 with a strong interest in research and industry in the intervening years, it would be logical to assume that edge computing would already be in widespread use today. Unfortunately, in practice, application developers rely only on the capabilities of the handheld device itself on the one end and the cloud on the other. There is usually nothing but the network between these two. To change the situation, there is a need for in-network computing resources, which should be deployed at the edge, preferably.

Initiatives from the industry, such as Multi-access Edge Computing (MEC), imply installing edge servers in the proximity of telecommunication masts. Also, cloud providers are addressing the issue by establishing new facilities closer to their users [2, 19, 22].

Although many of the above plans are likely to be realized in the near future, we see that crowdsourcing would complement these initiatives and enhance edge computing further. The main reason is the fact that computing devices at the disposal of the general public are exactly where they are needed the most — at the edge of the network. In the best possible case, clients would have a direct connection to them, avoiding the so-called last mile, often responsible for the significant fraction of network delay as recent measurement studies show [9, 23]. Also, in areas with the dense deployment of cloud datacenters, in-network edge would give mostly minor improvements [10].

Crowdsourcing has been a success driver in numerous initiatives and large-scale projects: Berkeley Open Infrastructure for Network Computing (BOINC) [1] and Folding@home [13], powering up decentralized finance [33]; enabling Peer-to-Peer (P2P) networks [4, 16, 18]. However, so far, crowdsourcing has not been playing any remarkable role in the development of edge computing, despite the obvious benefits it would bring. Although there are fully functional systems, they all suffer from poor crowd engagement: e.g., compare less than 1K nodes of Edge Network [12] to 700K machines in BOINC, providing 28.608 PetaFLOPS of computing power [5].

We see two major issues affecting the emergence of decentralized crowd-driven edge computing. First, the lack of a trusted coordinating platform, e.g., running networked services poses significant security threats for peers participating in general-purpose crowdsourced computing, involving legal responsibility in the worst-case scenarios. Second, there is no "big idea" backing edge computing, such as a search for extraterrestrial life, folding proteins, or creation of non-inflationary currency. In this position paper, we sketch multiple paths to overcome these issues or pose open questions for the community when there is no clear way forward.

The Big Three. The major cloud and client OS providers, Google, Microsoft, and Apple, have the potential to become game-changers for crowdsourced edge platforms. First, edge computing has always been intended to operate in strong coupling with the cloud [24], containing pre-stages of computation, so tight integration between the two is vital. Also, the reality is that the bulk of applications and services reside in the cloud today. Second, cloud and platform providers have control over the lion's share of end-user devices: over 70% of smartphones use the Android operating system developed by Google, the rest is dominated by Apple's IoS [29]. Of desktops, over 87% use Microsoft's Windows, and nearly 10% of the rest – Apple's Mac OS [25]. Google's Chrome OS is also gaining share. Third, cloud providers know the identities of their tenants and have legal agreements with them. Thus, the following chain of trust can be established: crowdsourcing peers – cloud providers – cloud tenants. Therefore, tenants can expand their applications to the edge via cloud providers, and the owners of the devices participating in the crowdsourcing can trust these applications since misbehaving tenants will be held legally responsible due to their agreements with the cloud providers. Additionally, there are virtualization and containerization technologies to protect end-user devices, as well as Trusted Execution Environments (TEEs) to guard tenants from malicious crowdsourcing peers. To summarize, participation of the general public in crowd-driven edge computing might be made as easy as selecting a checkbox in devices' preferences. Also, instead of cryptocurrencies, the rewards may come as credits that can be used in online stores of the *big three* – Microsoft, Google, Apple, and possibly other participants.

The scenario we presented above naturally has also downsides. Namely, the current state of the Internet is already characterized as "feudalism" [20], and there are notable appeals to revitalize it [3, 21]. Edge computing is seen as a unique opportunity to democratize the computation [26]. Quite obviously, what we suggest might appear only to exacerbate the situation. However, there is a positive side too. To avoid siloing and maximize the coverage of the edge services, the *big three* will need to cooperate so that, e.g., Google could leverage edge devices running Microsoft's Windows and vice versa. We believe this can potentially lead to the development of common standards and practices, eventually enabling the entrance of the third-party edge providers with cloudlets-like infrastructures, facilitating the emergence of dynamic edge environments [34].

Distributed Ledger Technologies (DLT). Barely anything has affected crowdsourcing in recent times more than the introduction of blockchain and its subsequent evolution into what is now called by the umbrella term DLT. Practically, every newly introduced crowdsourced system uses cryptographic tokens to attract peers, along with techniques like smart contracts to seal the agreements between parties and handle financial transactions securely.

The success of Ethereum [33], showing peoples' eagerness to let their machines run code of smart contracts, has inspired the development of a blockchain-driven system targeting general-purpose computation. A number of systems emerged, e.g., Golem [15], iExec [17], and SONM [28], to name a few, all gathering millions of initial funding in just a few hours from the enthusiastic public [7, 14, 32]. Nevertheless, despite the ability to attract investment, so far, most of the blockchain-driven systems have failed to engage decent numbers of participating peers. Possible reasons for the poor peer engagement can be the following. First, most of these proposals are targeting cloud computing. However, such crowdsourced-based systems cannot benefit from the economies of scale [31], which is the foundation of cloud computing's business success. Moreover, conventional cloud providers win also in technical performance and robustness: highly optimized network [11] interconnects cloud servers; also storage can be organized as SAN [30], etc. Second, general-purpose computation poses much higher security risks as compared to, e.g., Ethereums's smart contracts or pre-examined BOINC tasks, which are executed in isolated environments.

As the lack of economies of scale is inherent to edge computing due to its scattered deployment model, the problem is surmountable by creating systems targeting the edge specifically, like Edge Network [12]. Also, the above argument shows why drawbacks of crowdsourced cloud should not discourage crowd-driven edge. The elevated security risks are a more severe problem, and the solution here might be an inspection of deployable edge components by some trusted entities, such as Open Edge Computing Initiative [8].

Future Outlook. The scattered deployment model of edge computing offers a unique opportunity to decentralize the computation. Utilizing existing computing devices available in abundance at the disposal of general public, crowd-driven approaches would be more sustainable than establishing new datacenters. Both DLT and cloud providers can be the game-changers, removing the technical hindrances and security threats. The question remains what would be the "big idea" that would eventually make edge computing a socially remarkable phenomenon. Willing to support decentralization financially, people are not yet ready to let their hardware be utilized by novel platforms. We believe the combination of a viral killer application along with a technically mature platform will have the potential for the breakthrough of decentralized edge computing.

ACKNOWLEDGMENTS

We acknowledge the IWCI reviewers for useful feedback. This work was supported by the Swedish Foundation for Strategic Research with grant number GMT-14-0032 (Future Factories in the Cloud), the Academy of Finland in the BCDC (314167), AIDA (317086), and WMD (313477) projects.

REFERENCES

[1] Anderson, D. P. Boinc: a system for public-resource computing and storage. In *Fifth IEEE/ACM International Workshop on Grid Computing* (Nov 2004), pp. 4–10.
[2] AWS. Aws cloudfront. https://aws.amazon.com/cloudfront/, 2020.
[3] Balakrishnan, H., Banerjee, S., Cidon, I., Culler, D., Estrin, D., Katz-Bassett, E., Krishnamurthy, A., McCauley, M., McKeown, N., Panda, A., et al. Revitalizing the public internet by making it extensible. *ACM SIGCOMM Computer Communication Review 51*, 2 (2021), 18–24.
[4] BitTorrent. BitTorrent Protocol. https://www.bittorrent.com/, 2021.
[5] BOINC. https://boinc.berkeley.edu, 2020.
[6] Bonomi, F., Milito, R., Zhu, J., and Addepalli, S. Fog computing and its role in the internet of things. In *Proceedings of the first edition of the MCC workshop on Mobile cloud computing* (2012), ACM, pp. 13–16.
[7] Cointelegraph. SONM Ŝpectacularly SuccessfulЇCO Nets $42 Mln in Four Days. https://cointelegraph.com/news/sonm-spectacularly-successful-ico-nets-42-mln-in-four-days, 2020.
[8] Consortium. Open Edge Computing Initiative. https://www.openedgecomputing.org/, 2019.
[9] Corneo, L., Eder, M., Mohan, N., Zavodovski, A., and BayhanZ, S. Surrounded by the clouds. In *The Web Conference* (2021).
[10] Corneo, L., Mohan, N., Zavodovski, A., Wong, W., Rohner, C., Gunningberg, P., and Kangasharju, J. (how much) can edge computing change network latency? In *2021 IFIP Networking Conference (IFIP Networking)* (2021), pp. 1–9.
[11] Dalton, M., Schultz, D., Adriaens, J., Arefin, A., Gupta, A., Fahs, B., Rubinstein, D., Zermeno, E. C., Rubow, E., Docauer, J. A., et al. Andromeda: Performance, isolation, and velocity at scale in cloud network virtualization. In *15th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 18)* (2018), pp. 373–387.
[12] Edge Network. https://edge.network, 2020.
[13] Foldinghome. https://foldingathome.org, 2020.
[14] Forbes, Roger Aitken. Fintech golem's 'airbnb' for computing crowdsale scores $8.6m in minutes. https://www.forbes.com/sites/rogeraitken/2016/11/12/fintech-golems-airbnb-for-computing-crowdsale-scores-8-6m-in-minutes, November 2016.
[15] Golem Worldwide Supercomputer. https://golem.network/, 2020.
[16] I2P. The Invisible Internet Project. https://geti2p.net/, 2021.
[17] iExec. https://iex.ec/, 2020.
[18] InterPlanetary File System. IPFS powers the Distributed Web. https://ipfs.io/, 2021.
[19] Kurian, T. Google cloud unveils strategy for telecommunications industry. https://cloud.google.com/blog/topics/inside-google-cloud/google-cloud-unveils-strategy-telecommunications-industry, 2020.
[20] Liu, T., Tariq, Z., Chen, J., and Raghavan, B. The barriers to overthrowing internet feudalism. In *Proceedings of the 16th ACM Workshop on Hot Topics in Networks* (2017), pp. 72–79.
[21] McCauley, J., Harchol, Y., Panda, A., Raghavan, B., and Shenker, S. Enabling a permanent revolution in internet architecture. In *Proceedings of the ACM Special Interest Group on Data Communication* (New York, NY, USA, 2019), SIGCOMM '19, Association for Computing Machinery, p. 1–14.
[22] Microsoft. Azure regions. https://azure.microsoft.com/en-us/global-infrastructure/regions/, 2020.
[23] Mohan, N., Corneo, L., Zavodovski, A., Bayhan, S., Wong, W., and Kangasharju, J. Pruning edge research with latency shears. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks* (2020), pp. 182–189.
[24] Mohan, N., and Kangasharju, J. Edge-fog cloud: A distributed cloud for internet of things computations. In *2016 Cloudification of the Internet of Things (CIoT)* (Nov 2016), pp. 1–6.
[25] Netmarketshare. Operating System Market Share. https://netmarketshare.com/operating-system-market-share.aspx, 2021.
[26] Peterson, L., Anderson, T., Katti, S., McKeown, N., Parulkar, G., Rexford, J., Satyanarayanan, M., Sunay, O., and Vahdat, A. Democratizing the network edge. *ACM SIGCOMM Computer Communication Review 49*, 2 (2019), 31–36.
[27] Satyanarayanan, M., Bahl, P., Caceres, R., and Davies, N. The case for vm-based cloudlets in mobile computing. *IEEE pervasive Computing*, 4 (2009), 14–23.
[28] Sonm. https://sonm.com/, 2020.
[29] StatCounter. Mobile Operating System Market Share Worldwide. https://gs.statcounter.com/os-market-share/mobile/worldwide/, 2021.
[30] Tate, J., Beck, P., Ibarra, H. H., Kumaravel, S., Miklas, L., et al. *Introduction to storage area networks*. IBM Redbooks, 2018.
[31] Weinman, J. *Cloudonomics: The business value of cloud computing*. John Wiley & Sons, 2012.
[32] William Suberg. iEx.ec Closes World's 5th Largest ICO With $12 mln in 6 Hours. https://cointelegraph.com/news/iexec-closes-worlds-5th-largest-ico-with-12-mln-in-6-hours, 2017.
[33] Wood, G. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum Project Yellow Paper* (2014).
[34] Zavodovski, A., Mohan, N., Bayhan, S., Wong, W., and Kangasharju, J. ICON: Intelligent container overlays. In *Proceedings of the 17th ACM Workshop on Hot Topics in Networks* (New York, NY, USA, 2018), HotNets '18, ACM, pp. 15–21.