Decentralised Edge-Computing and IoT through Distributed Trust

Ioannis Psaras
University College London
i.psaras@ucl.ac.uk

ABSTRACT

The emerging Internet of Things needs edge-computing - this is an established fact. In turn, *edge computing needs infrastructure decentralisation*. What is not necessarily established yet is that *infrastructure decentralisation needs a distributed model of Internet governance and decentralised trust schemes*. We discuss the features of a decentralised IoT and edge-computing ecosystem and list the components that need to be designed, as well the challenges that need to be addressed.

CCS CONCEPTS

• **Security and privacy → Network security**; • **Networks → Network protocol design**; **Routing protocols**; **Network components**; *Naming and addressing*;

KEYWORDS

Edge-Computing, Distributed Trust, Programmable Privacy, Blockchain

ACM Reference Format:
Ioannis Psaras. 2018. Decentralised Edge-Computing and IoT through Distributed Trust. In *MobiSys '18: The 16th Annual International Conference on Mobile Systems, Applications, and Services, June 10–15, 2018, Munich, Germany*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3210240.3226062

1 COMPUTE DECENTRALISATION

During the past twenty years (or so) we have been witnessing a continuous trend towards centralising Internet content delivery and application-oriented computation. Centralisation led to the development of massive scale data-centres (commonly referred to as "the cloud"), which is the place where 90% of user requests end up being executed. Although this trend served well the purpose of the Internet as we know it today, and was also inline with the demand of economies of scale, it is certainly not fit for purpose for future applications. The 5G architecture and the emerging Internet of Things will demand for applications that respond in sub-msec latencies. Such applications cannot tolerate the response-latency of centralised computations siloed behind closed walls, typically in far-away data-centres.

Edge-/Fog-computing has been proposed as a complementary paradigm to the cloud. Its main premise is the de-centralisation of the cloud into multiple smaller scale computing devices, or cloudlets (ranging from mini-data centres to WiFi APs and to single-board computers), largely referred to as "computation spots".

The expectation from the mobile edge-/fog-computing paradigm presents, to a certain extent, similarities to the caching era of the 90s. That is, similarly to the move from servers acting as the sole providers of static content to proxy caches, and more recently ubiquitous in-network caches (in the ICN area) [34], [33], [19], [35], [24] the edge-computing paradigm is attempting a shift of computation closer to the users [13], [31], [11], [25], [14]. By and large, the rationale behind deploying proxy caches was: *i)* reduce response delay to end-users, *ii)* reduce core-network traffic, and, *iii)* reduce server load. Moving to the edge-/fog-computing paradigm, we could realistically argue that the motivation and expectation is roughly similar: move network functions and user-facing applications closer to the users to reduce response delay and network traffic towards the core [13], [10], [27], [23].

Interestingly, there is one more dimension that can severely affect the performance of an edge-computing system - as opposed to proxy-caching functionality. That is, resolving functions, *i.e.*, computation functionality, on the fly is impossible to handle by the current DNS infrastructure. Functions can get instantiated and dissolved in a matter of seconds and need to be resolved and executed within *msecs*, while normal DNS entries are updated a few orders of magnitude slower, *i.e.*, in the order of minutes, if not much longer. A computation-centric paradigm, where functions are packaged in lightweight virtualisation environments [26], *e.g.*, Unikernels [21], [22], and are explicitly addressed in user requests is a central component of the system. Requests carry input parameters for the function, while functions are stateless and can therefore be executed at any network node (see Serverless architecture [15]).

2 REMOVING TRUST

There is one extra challenge in case of distributed edge and fog computing, which has not been encountered before and which is likely to influence significantly adoption of the paradigm: **Internet infrastructure governance**, or in order words, *who is owning and managing the edge-computing infrastructure and who to trust when using edge Internet services*. In the current Internet landscape, infrastructure is owned and operated (in obscure ways) by an oligopoly of "tech giants" - the likes of Google, Microsoft, Amazon, Facebook and Akamai. This model has worked relatively well in terms of performance[1] so far, but it is questionable whether a similar model would work well in case of a distributed edge computing environment.

Firstly, it is easy to technically manage (and provide relatively acceptable performance for) a few centralised computation factories, but almost impossible to manage and administer billions of computation spots centrally. Secondly, innovation reaches a threshold difficult to pass when infrastructure stays behind closed doors.

[1] Although it has failed its users hugely in terms of privacy.

MobiSys '18, June 10–15, 2018, Munich, Germany

ACM ISBN 978-1-4503-5720-3/18/06.
https://doi.org/10.1145/3210240.3226062

Thirdly, it is embarrassing to witness that after 40 years of intense research, engineering and development, if a link to the centralised infrastructure fails the most basic Internet functionalities break.[2] Instead, it is reasonable to assume that a Distributed Edge Computing infrastructure, which will be responsible for vital applications, such as driving our cars, will be run by a multitude of players operating closer to the end-user. Clearly, performance, security and privacy need to come to the forefront of attention, as such applications will soon be responsible for extremely latency-sensitive applications that will manage central and in many cases life-threatening aspects of our lives.

To remove centralisation is to remove the trust from the infrastructure provider.

Whether one trusts the tech giants or not, by using their infrastructure one silently accepts that they will do their best to provide high performance, security and privacy. Moving to a decentralised and distributed governance model, users will have to trust unknown operators, companies and platforms (mostly in the form of Decentralised Autonomous Organisations - DAOs [32]), effectively removing trust from today's Internet ecosystem landscape.

Recent advances in cryptography and Distributed Ledger Technology (*aka* blockchain) can be of significant contribution at this point. Distributed ledgers can track and record any transaction between any two entities in a trustless manner in an immutable history record. Security can be improved and privacy can be guaranteed. Despite performance issues of current blockchain systems [8], [20] (which are receiving significant attention and are expected to be solved in the near future), the important point is this:

Computing infrastructure can be distributed to billions of storage and computation spots, operated by anyone who can innovate on it, while governance can become decentralised and guarantee higher levels of security and privacy than the current infrastructure.

Distributed computing between trustless nodes is an enabler for ubiquitous computing, where any spare computation cycle or memory space can be exploited to store and execute latency-sensitive applications in geographically close locations. In turn, latency to reach the computation spot is reduced, execution time within the computation spot is kept to a minimum and applications are guaranteed to receive timely responses.

3 COMPONENTS AND CHALLENGES

Computation-Centric Architecture: New computation-centric architectures are needed to address the need for fast resolution of network functions that are executed at the edges of the network [17], [28]. The ultimate purpose of such architectures is to alleviate the need for costly, in terms of latency, communication to DNS-like resolution services and therefore, improve end-user and application QoS. In order to address this challenge, we need to move away from the current host-based paradigm and adopt an Information- and Computation-Centric model, where user requests are directly addressing Information [34] and Computation [17].

[2] Amazon Web Services (AWS) holds a 40% share of the cloud-server market. When AWS's Virginia datacenter had an outage, a significant part of the web went offline [29].

Secure Payment System: In an open, non-walled garden and decentralised cloud computing environment, execution nodes are owned by multiple stakeholders, while requestors do not know which nodes will execute their tasks and thus whom to pay. The threat model has several dimensions: users do not want to pay for yet unfinished or unverified tasks; on the other hand, an execution node receiving a request does not want to use its resources without making sure that it will eventually be paid. An efficient and secure payment system is essential and can determine the future success of decentralised storage and computing [9], [18]. Accounting, accountability and transaction verification can be supported through blockchain technology, while a robust design for micropayments can be implemented on top of off-chain payment channels [9].

Rewards: The peer-to-peer networking area was hugely successful, but has failed to get adopted as the mainstream model for content distribution. Although part of the reason was the fact that ISPs were not comfortable with the traffic model, limited user participation (*i.e.*, seeding) played a significant role. It turns out that *rewards* are a very important part of any decentralised system. That said, moving to a decentralised Internet infrastructure where any individual or company can provide own resources for content storage/delivery and computation needs to be coupled with a robust, fair and scalable business model and reward scheme.

Several platforms are being developed at the moment to decentralise storage [4], [7], [5] and computing [2], [6], [3], [1] and reward end-users for their contribution to the network. Reward models based on cryptographic tokens (*aka* cryptocurrencies) have not been developed yet and the new area of *cryptoeconomics or tokenomics* is emerging as an essential component of the system. Cryptoeconomic models need to take into account a variety of factors ranging from technical and objective factors ([14], [16], [30]), *e.g.*, system stability and scalability, resource scarcity and convergence to equilibrium, to less-technical and subjective ones, *e.g.*, user-perceived privacy, QoE and incentives.

Programmable Privacy: Privacy-violating services and security breaches from cloud-based IoT companies are revealed almost daily. This has made users very reluctant to engage to new services and is likely to influence the overall adoption and integration of edge-computing in general and IoT in particular. Identity verification and access control to private user data ([12]) can be automatically granted (or denied) if specific rules, programmed in smart contracts (and recorded on the blockchain), are met. The community needs to not only re-define user-privacy, but also develop new software structures to grant access to sensitive personal data if the rules programmed in the smart contract are met, but most importantly support access right revocation if specific conditions change over time (*e.g.*, if personal data get used in undesired ways).

The ultimate challenge is therefore, to bridge the gap between the established Internet infrastructure and related protocols and the new development activities in the areas of information security, privacy and distributed ledger technology. Instead of being seen as an afterthought, security, privacy and trust need to come to the forefront of research, design and development. The amalgamation of these areas at large can together lay the ground for the deployment of decentralised, distributed, trustless edge computing and ultimately the Internet of Things.

REFERENCES

[1] [n. d.]. Blockstack: The New Decentralised Internet. ([n. d.]). https://blockstack.org/

[2] [n. d.]. Golem Network. ([n. d.]). https://golem.network/

[3] [n. d.]. iExec Network: Blockchain-Based Decentralized Cloud Computing. ([n. d.]). https://iex.ec/

[4] [n. d.]. IPFS: Interplanetary File System. ([n. d.]). https://ipfs.io

[5] [n. d.]. MaidSafe Network. ([n. d.]). https://maidsafe.net/

[6] [n. d.]. SONM: Decentralized Fog Computing Platform. ([n. d.]). https://sonm.com/

[7] [n. d.]. StorJ: Distributed Cloud Storage. ([n. d.]). https://storj.io

[8] Mustafa Al-Bassam, Alberto Sonnino, Shehar Bano, Dave Hrycyszyn, and George Danezis. 2018. Chainspace: A Sharded Smart Contracts Platform. In *In Proceedings of the Network and Distributed System Security Symposium (NDSS)*.

[9] Mustafa Al-Bassam, Alberto Sonnino, Michal Król, and Ioannis Psaras. 2018. Airtnt: Fair Exchange Payment for Outsourced Secure Enclave Computations. (2018).

[10] O. Ascigil, T. K. Phan, A. G. Tasiopoulos, V. Sourlas, I. Psaras, and G. Pavlou. 2017. On Uncoordinated Service Placement in Edge-Clouds. In *2017 IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*. 41–48. https://doi.org/10.1109/CloudCom.2017.46

[11] E. Bastug, M. Bennis, and M. Debbah. 2014. Living on the edge: The role of proactive caching in 5G wireless networks. *IEEE Communications Magazine* 52, 8 (Aug 2014), 82–89. https://doi.org/10.1109/MCOM.2014.6871674

[12] Amir Chaudhry, Jon Crowcroft, Heidi Howard, Anil Madhavapeddy, Richard Mortier, Hamed Haddadi, and Derek McAuley. 2015. Personal Data: Thinking Inside the Box. In *Proceedings of The Fifth Decennial Aarhus Conference on Critical Alternatives (AA '15)*. Aarhus University Press, 29–32. https://doi.org/10.7146/aahcc.v1i1.21312

[13] Sarah Clinch et al. 2012. How close is close enough? Understanding the role of cloudlets in supporting display appropriation by mobile users. In *PerCom*. IEEE.

[14] Patricia Takako Endo et al. 2011. Resource allocation for distributed cloud: concepts and research challenges. *IEEE Network* 25, 4 (2011).

[15] Scott Hendrickson, Stephen Sturdevant, Tyler Harter, Venkateshwaran Venkataramani, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. 2016. Serverless Computation with openLambda. In *Proceedings of the 8th USENIX Conference on Hot Topics in Cloud Computing (HotCloud'16)*. USENIX Association, Berkeley, CA, USA, 33–39. http://dl.acm.org/citation.cfm?id=3027041.3027047

[16] A-Long Jin et al. 2016. Auction mechanisms toward efficient resource sharing for cloudlets in mobile cloud computing. *Transactions on Services Computing* (2016).

[17] Michal Krol and Ioannis Psaras. 2017. NFaaS: Named Function as a Service. In *ACM ICN'17*. ACM, 1–11.

[18] Michal Krol and I. Psaras. 2018. Secure Payments for Outsourced Computations. In *2018 NDSS Workshop on Decentralised IoT Security and Standards*.

[19] D. Kutscher and et al. [n. d.]. RFC 7927: Information-Centric Networking (ICN) Research Challenges. ([n. d.]). https://tools.ietf.org/html/rfc7927

[20] Joshua Lind, Ittay Eyal, Peter R. Pietzuch, and Emin Gün Sirer. 2016. Teechan: Payment Channels Using Trusted Execution Environments. *CoRR* abs/1612.07766 (2016). arXiv:1612.07766 http://arxiv.org/abs/1612.07766

[21] Anil Madhavapeddy, Thomas Leonard, Magnus Skjegstad, Thomas Gazagnaire, David Sheets, et al. [n. d.]. Jitsu: Just-In-Time Summoning of Unikernels.. In *NSDI, 2015*.

[22] Anil Madhavapeddy and David J. Scott. 2013. Unikernels: Rise of the Virtual Library Operating System. *Queue* 11, 11 (2013).

[23] Ioannis Psaras, Onur Ascigil, Sergi Rene, George Pavlou, Alex Afanasyev, and Lixia Zhang. 2018. Mobile Data Repositories at the Edge. In *Proceedings of the 1st USENIX Workshop on Hot Topics in Edge Computing (HotEdge'18)*.

[24] I. Psaras, W. K. Chai, and G. Pavlou. 2014. In-Network Cache Management and Resource Allocation for Information-Centric Networks. *IEEE Transactions on Parallel and Distributed Systems* 25, 11 (Nov 2014), 2920–2931. https://doi.org/10.1109/TPDS.2013.304

[25] C. A. Sarros, S. Diamantopoulos, S. Rene, I. Psaras, A. Lertsinsrubtavee, C. Molina-Jimenez, P. Mendes, R. Sofia, A. Sathiaseelan, G. Pavlou, J. Crowcroft, and V. Tsaoussidis. 2018. Connecting the Edges: A Universal, Mobile-Centric, and Opportunistic Communications Architecture. *IEEE Communications Magazine* 56, 2 (Feb 2018), 136–143. https://doi.org/10.1109/MCOM.2018.1700325

[26] Mahadev Satyanarayanan et al. 2009. The case for VM-based cloudlets in mobile computing. *Pervasive Computing* (2009).

[27] E. M. Schooler, D. Zage, J. Sedayao, H. Moustafa, A. Brown, and M. Ambrosin. 2017. An Architectural Vision for a Data-Centric IoT: Rethinking Things, Trust and Clouds. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. 1717–1728. https://doi.org/10.1109/ICDCS.2017.243

[28] Manolis Sifalakis, Basil Kohler, Christopher Scherb, and Christian Tschudin. 2014. An information centric network for computing the distribution of computations. In *ACM ICN'14*. 137–146.

[29] Jake Swearingen. 2018. http://nymag.com/selectall/2018/03/when-amazon-web-services-goes-down-so-does-a-lot-of-the-web.html. (2018).

[30] Argyrios G. Tasiopoulos, Onur Ascigil, Ioannis Psaras, and George Pavlou. 2018. Edge-MAP: Auction Markets for Edge Resource Provisioning. In *WoWMoM'18*. IEEE.

[31] Tim Verbelen et al. 2012. Cloudlets: Bringing the cloud to the mobile user. In *Mobile cloud computing and services workshop*. ACM.

[32] Wikipedia. [n. d.]. Distributed Autonomous Organisations. ([n. d.]). https://en.wikipedia.org/wiki/Decentralized_autonomous_organization

[33] G. Xylomenos, C. N. Ververidis, V. A. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K. V. Katsaros, and G. C. Polyzos. 2014. A Survey of Information-Centric Networking Research. *IEEE Communications Surveys Tutorials* 16, 2 (Second 2014), 1024–1049. https://doi.org/10.1109/SURV.2013.070813.00063

[34] Lixia Zhang, Alexander Afanasyev, Jeffrey Burke, Van Jacobson, kc claffy, Patrick Crowley, Christos Papadopoulos, Lan Wang, and Beichuan Zhang. 2014. Named Data Networking. *SIGCOMM Comput. Commun. Rev.* 44, 3 (July 2014), 66–73. https://doi.org/10.1145/2656877.2656887

[35] M. Zhang, H. Luo, and H. Zhang. 2015. A Survey of Caching Mechanisms in Information-Centric Networking. *IEEE Communications Surveys Tutorials* 17, 3 (thirdquarter 2015), 1473–1499. https://doi.org/10.1109/COMST.2015.2420097