

To Edge or Not to Edge?

Faria Kalim, Shadi A. Noghbi, Shiv Verma

University of Illinois at Urbana-Champaign, [kalim2, abdolla2, sverma11]@illinois.edu

ABSTRACT

Edge computing caters to a wide range of use cases from latency sensitive to bandwidth constrained applications. However, the exact specifications of the edge that give the most benefit for each type of application are still unclear. We investigate the concrete conditions when the edge is feasible, i.e., when users observe performance gains from the edge while costs remain low for the providers, for an application that requires both low latency and high bandwidth: video analytics.

KEYWORDS Edge Computing, Distributed Systems

Introduction: Edge computing pushes computation from the cloud to the edge of the network. The concept has recently gained tremendous momentum in both academia and industry as it can enable a wide range of use cases from latency sensitive to bandwidth/energy constrained, and privacy concerned applications [2, 4]. Examples include video analytics, online gaming, virtual reality, and self-driving cars.

Despite the clear potential benefits of the edge, it is unclear what exact specifications the edge should have for its benefits to bloom. This is particularly important for capacity planning (determining how much to invest to see performance gains) and feasibility verification (checking whether the current hardware can meet requirements). Here, we investigate the concrete conditions when an edge is *feasible*, i.e., when users are able to observe performance gains from the edge, at low cost for the service provider. Specifically, we try to answer three main questions: 1) Since locality is key for edge computing, how close should the edge be to the client applications? 2) Do we require as much bandwidth to the edge as we need to the cloud? 3) How much increased load can the edge sustain?

The Testbed: We study the utility of the edge in video analytics applications such as those that analyze traffic patterns near an intersection to warn self-driving cars of passing pedestrians. Video analytics is usually both bandwidth heavy (generating massive data volumes) and latency sensitive (e.g., acceptable latency is below 185 ms for video in humans [5] and a few tens of ms for self-driving cars), making the edge a great fit for it.

Our testbed allows us to vary the key factors that are relevant to our application, i.e., latency, bandwidth and the power of the edge. Thus we can mimic various environments, measure throughput and latency in each case, and determine when using the edge achieves significant performance gains.

We built an Android application that captures pictures and performs object detection on them. We offload object detection to a remote server (edge or cloud) that trains a neural network

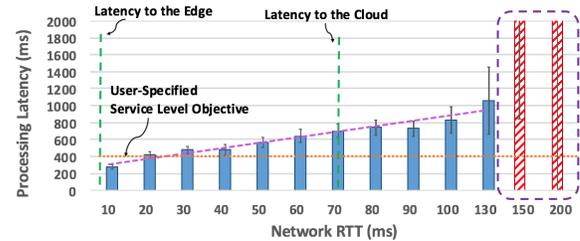


Figure 1: Latency for a processed request increases linearly as network RTT increases. After a threshold, this can cause a build-up of requests at the sender, causing unbounded latencies (shown by red bars).

based model using the YOLO library [3], employs it to detect objects in received pictures, and returns them to the client.

Evaluation and Takeaways: We present one evaluation result in Fig. 1, which shows that even less powerful machines can be used as edge-servers, if the latency to the edge is small. We increase network latency from the client to the edge using Netem. We have a 1 Gbps connection, with an edge provisioned with an NVIDIA Titan X Graphics Card for object detection. Intuitively, the likelihood that latency SLOs will be violated increases in proportion to network RTT. As latencies increase, requests begin to build up at the sender, causing unbounded response latencies. This impacts the cloud more than the edge as it is not only farther away from clients but also has incoming requests from more clients than the edge.

Our takeaways are: edges can be treated as a resource for a limited number of users who are closest to it. Since this means that there will be a lot less load on the edge, it is possible for resource providers to provision less powerful machines and less bandwidth to the edge. However, they need to ensure that sufficient edges are provisioned such that the network latency to the user is minimized. Additionally, the benefits of the edge truly become obvious when the computation time (for both the edge and the cloud) is less than the network latency. However, an increase in load affects the edge the same way as the cloud: the network or compute resources can become a bottleneck, causing unbounded latencies. The exact specification of machines and bandwidth vary from application to application. For instance in our setting, an edge placed at 20 ms is ideal.

In future work, we will consider building on recent work [1] to construct machine learning models that help determine the edge specifications required for a certain number of users and type of workload, while minimizing cost of infrastructure.

REFERENCES

- [1] Omid Alipourfard, Hongqiang Liu, Jianshu Chen, Shivaram Venkataraman, Minlan Yu, and Ming Zhang. 2017. CherryPick: Adaptively Unearthing the Best Cloud Configurations for Big Data Analytics. In *NSDI*.
- [2] Flavio Bonomi, Rodolfo Milito, Jiang Zhu, and Sateesh Addepalli. 2012. Fog computing and its role in the internet of things. In *Proceedings of the first edition of the MCC workshop on Mobile cloud computing*. ACM, 13–16.
- [3] Joseph Redmon. 2013–2016. Darknet: Open Source Neural Networks in C. <http://pjreddie.com/darknet/>. (2013–2016).
- [4] Mahadev Satyanarayanan, Pieter Simoons, Yu Xiao, Padmanabhan Pillai, Zhuo Chen, Kiryong Ha, Wenlu Hu, and Brandon Amos. 2015. Edge analytics in the internet of things. *IEEE Pervasive Computing* 14, 2 (2015).
- [5] Wikipedia. 2017. Latency (audio) — Wikipedia, The Free Encyclopedia. (2017). [https://en.wikipedia.org/w/index.php?title=Latency_\(audio\)&oldid=790236389](https://en.wikipedia.org/w/index.php?title=Latency_(audio)&oldid=790236389) [Online; accessed 13-July-2017].

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). SoCC '17, September 24–27, 2017, Santa Clara, CA, USA

© 2017 Copyright held by Owner/Author
ACM ISBN 978-1-4503-5028-0/17/09...\$15.00
DOI: <https://doi.org/10.1145/3127479.3132572>