PhD Forum: Crowd-sourcing Mobile Devices to Provide Storage in Edge-Clouds

João A. Silva, João M. Lourenço, and Hervé Paulino*

NOVA LINCS/Departamento de Informática, Faculdade de Ciências e Tecnologia
Universidade NOVA de Lisboa, 2829-516 Caparica, Portugal
jaa.silva@campus.fct.unl.pt {joao.lourenco,herve.paulino}@fct.unl.pt

Abstract. Given the proliferation and enhanced capabilities of mobile devices, their computational and storage resources can now be combined in a wireless cloud of nearby mobile devices, a mobile edge-cloud. These clouds are of particular interest in low connectivity scenarios, e.g., sporting events and disaster scenarios. We are particularly interested in supporting storage services in these kind of edge-clouds, as a mean to enable data sharing, dissemination and querying. In this PhD thesis, we propose to address this concern by researching on the usage of ad-hoc clouds of mobile devices to develop an efficient storage service capable of providing high availability and reliability.

1 Introduction

Since the advent of mobile computing, mobile applications have been developed under the assumption that mobile devices are *thin clients*: resource-limited devices that have to rely on backend (thick) servers to do all the heavy computations. However, with the rapid growth of these devices' capabilities, we may currently find mobile devices with more computational and storage power than the desktops of the last decade [2]. This makes possible to start seeing mobile devices as *thick clients*, or even as *thin servers*. Another interesting aspect is the extensive proliferation of mobile devices, nearly *7 billion* worldwide, a number bound to increase with the current demand in developing countries. We are thus witnessing a significant growth in the number and density of these devices.

Given the proliferation and enhanced capabilities of today's mobile devices, it is now a real possibility to consider the opportunities that a *collection of mobile devices* might *collaboratively* present as a computational and storage resource.

Recently, Drolia et al. introduced the notion of *mobile edge-cloud* [5]—an *opportunistic* cloud made entirely of mobile devices with *high mobility* in *close geographical proximity* that work as its computational and storage nodes. Essentially, *all* of the cloud's nodes are edge (and not server-calibre) devices, and *all* computations are performed completely within the edge-cloud. These clouds are appropriate in scenarios where high bandwidth, low latency connections to the Internet may be limited, e.g., in remote locations or in disaster response scenarios, and also when infrastructure connectivity is available, but the quality of service is degraded beyond use due to congestion.

* This work is partially funded by FCT-MEC in the context of the research project CMUP-ERI/FIA/0048/2013.

Cloud infrastructures subsume mechanisms to efficiently handle resource discovery and utilization, by relying on centralized services that cannot be easily ported into a fully decentralized and dynamic environment such as mobile edge-clouds. As such, in this PhD thesis, we are targeting storage services for mobile edge-clouds, and we will consider many challenges, including: (i) support heterogeneous edge devices; (ii) efficiently detect and tolerate churn and network partitions; (iii) provide low latency operations (e.g., efficient location of information); (iv) minimize energy consumption; (v) ensure data consistency; (vi) endure poor/zero connectivity; and (vii) provide transactional support.

With this PhD thesis we plan to develop a storage service for mobile edge-clouds, akin from the ones provided by standard infrastructure clouds. We propose to investigate how to leverage both strong and eventual consistency techniques to develop resource-aware algorithms for dealing with churn, high network dynamism and device heterogeneity, among other aspects. Our main expected contributions are: (i) a set of algorithms leveraging on both strong and eventual consistency techniques to provide distributed storage for edge-clouds, tolerating churn and network partitions; (ii) a prototype of a storage solution for edge-clouds that incorporates the proposed algorithms; and (iii) a detailed analysis of the trade-offs between strong and eventual consistency as applied in our prototyped storage solution, obtained from both simulation and real world use cases.

2 State of the Art

Hyrax [9] is a port of the Hadoop MapReduce implementation to Android. Other works have also developed mobile MapReduce implementations [4] but they only provide distributed computation, whereas Hyrax also provides storage. Moreover, they are not completely mobile, and do not offer measures to overcome inherent obstacles in mobile environments, e.g., churn and resource constraints. Other works [3] also suggest computation offload, in order to conserve power on mobile devices. They propose mechanisms to augment resource-constrained devices with non-mobile infrastructures. Conversely, the idea of mobile edge-clouds proposes to build ad-hoc clouds and provide a novel collaborative execution environment.

Distributed storage systems such as OceanStore [6] use distributed hash tables (DHTs) as a building block to manage the mapping from a collection of keys to values without any centralized control. DHTs are designed to handle node arrivals, departures, and failures gracefully. Edge-clouds also have to tolerate these events, but the highly volatile nature of the underlying network raises new challenges in the provision of fast lookups.

PAN [8] provides reliable storage in mobile ad-hoc networks (MANETs), and it uses probabilistic quorum techniques to deal with dynamic and unpredictable topology changes. The main goal is to provide high availability to small data objects, while supporting high reliability under different access and update rates.

Similarly to Hyrax, [1] introduces the idea of using distributed voluntary resources to form dispersed, less-managed clouds—nebulas. The notion of edge-cloud argues that we can harness the resources of everyday mobile devices, while nebulas aim at leveraging end-user hosts such as workstations.

Phoenix [10] is the work that more closely relates to what we are aiming for. The authors propose a storage system using an autonomous mobile infrastructure, but its primary goal is to maintain the desired level of data storage redundancy fairly accurate, and does not support data updates directly. Conversely, storage in edge-clouds aims at being a more generic storage system for a network of common, everyday mobile devices.

3 Approach

We plan to leverage on peer-to-peer (P2P) techniques and others such as adaptive replication (AdR), data structures with reconciliation (e.g., commutative replicated data types (CRDTs) [11]), and the combination of both strong and eventual consistency (e.g., RedBlue consistency [7]). Preguiça et al. [11] present a principled approach to build systems supporting weak consistency using data types called CRDTs, which ensure the convergence of all replicas irrespectively of the order by which updates are applied. Li et al. introduced RedBlue consistency [7]—a way of building geo-replicated systems that use weak consistency whenever possible but resort to strong consistency whenever necessary.

During the development of this work, we will address four main requirements:

Availability. To tackle network partitions and poor connectivity (challenges (ii) and (vi)) we will explore AdR, CRDTs, and the combination of strong and eventual consistency. Since CRDT operations can commute with each other, they allow devices to be disconnected or partitioned in the network and still be able to work. When the connection is re-established, devices can reconcile and merge their state, converging to a common final state. Regarding AdR, if data hot-spots have a higher replication factor, the system is able to serve more requests for those items, while occasionally accessed items just need to be replicated according to the minimum requirements. The combination of both strong and eventual consistency will also allow the system to abide network partitions and still continue to operate (albeit in a more limited way) despite of device disconnections. Challenge (ii) also mentions the problem of churn, which P2P techniques, as well as AdR, can help solving.

Scalability. The unstructured overlay topologies of P2P networks can be used to tackle challenge (iii)—provide low latency operations—by maintaining the membership of devices in the edge-clouds, as well as by efficiently propagating information and locating relevant resources. As the edge-clouds grow, being able to use eventual consistency can become a very big advantage, since it requires minimal coordination among replicas.

Usability. Eventual consistency is known for being harder to reason with, when comparing to strong consistency models. Our combination of strong and eventual consistency shall be established in a familiar way to the programmer, in order to ease the burden of reasoning with it (challenge (v)). Due to the evolution of mobile devices' capabilities, their main resource constraints are bandwidth and energy (challenge (iv)). For a system like this to be adopted, it must be both energy and bandwidth efficient, which can be achieved through load balancing, and by maximizing local operations while minimizing communication. The use of transactions (with ACID-like semantics) to maintain consistency when accessing data (challenge (vii)) can also be advantageous, since it is largely used by distributed storage systems, and is very familiar to programmers.

Resource-Awareness. Mobile devices are still resource-constrained, as such the developed techniques and algorithms shall have into account resource-awareness. To minimize energy consumption data needs to be replicated, and data replicas need to be accessed and/or modified with minimal coordination. Challenge (i) is intrinsically a system's aspect that can be dealt with using load balancing techniques taking into account the available resources.

Evaluation. All the developed algorithms and techniques will be validated according to multiple criteria, e.g., efficiency, usability, and scalability. We will develop specific applications to integrate our solution into, and devise benchmarks enabling an empirical evaluation. We aim to evaluate our storage solution under a variety of conditions, e.g., poor/zero connectivity, different degrees of churn, and different device densities. The developed applications will be evaluated in both simulation and real world environments.

4 Conclusions

In this PhD thesis, we will develop algorithms and system solutions to build a storage service for mobile edge-clouds. This storage system must be able to provide highly available and reliable storage in clouds of mobile devices without resorting (or as little as possible) to infrastructure clouds. Our main focus will be on leveraging on the combination of both strong and eventual consistency, and its integration with CRDTs to develop resource-aware algorithms for dealing with churn, high network dynamism and devices' heterogeneity.

References

1. A. Chandra et al. Nebulas: Using distributed voluntary resources to build clouds. In *HotCloud*, 2009.
2. H.-D. Cho et al. Benefits of the big.LITTLE Architecture. Technical report, Samsung, Feb 2012.
3. B.-G. Chun et al. Clonecloud: Elastic execution between mobile device and cloud. In *EuroSys*, 2011.
4. D. Diaz-Sanchez et al. Flexible computing for personal electronic devices. In *ICCE*, 2013.
5. U. Drolia et al. The case for mobile edge-clouds. In *UIC/ATC*, 2013.
6. J. Kubiatowicz et al. Oceanstore: An architecture for global-scale persistent storage. *SIGPLAN Not.*, 2000.
7. C. Li et al. Making geo-replicated systems fast as possible, consistent when necessary. In *OSDI*, 2012.
8. J. Luo et al. Pan: Providing reliable storage in mobile ad hoc networks with probabilistic quorum systems. In *MobiHoc*, 2003.
9. E. Marinelli. Hyrax: Cloud computing on mobile devices using mapreduce. Master's thesis, Carnegie Mellon University, 2009.
10. R. Panta et al. Phoenix: Storage using an autonomous mobile infrastructure. *TPDS*, 2013.
11. N. Preguiça et al. A commutative replicated data type for cooperative editing. In *ICDCS*, 2009.