Decentralized Storage for Networks of Hand-held Devices

Ricardo Monteiro João Silva João Lourenço Hervé Paulino
NOVA Laboratory for Computer Science and Informatics & Departamento de Informática
Faculdade de Ciências e Tecnologia, Universidade Nova de Lisboa, 2829-516 Caparica, Portugal
{rad.monteiro,jaa.silva}@campus.fct.unl.pt {joao.lourenco,herve.paulino}@fct.unl.pt

ABSTRACT

In this paper we propose a fully distributed storage system for everyday hand-held mobile devices, e.g., smartphones and tablets, that follows a best effort approach to ensure data persistence and availability even in the presence of churn (i.e., the unpredictable arrival and departure of nodes).

Keywords

Distributed Storage, Mobile Networks, Peer-to-Peer.

1. INTRODUCTION

The power of today's mobile devices can be compared to the one of desktops a few years old. This recent increase in the computational and storage capabilities (of mobile devices) enables their use in much more complex tasks. An example is the formation of a cloud via the aggregation of the individual resources of a set of mobile devices, in order to share and process data generated within such cloud [2]. For instance, a collection of 8 smartphones, each one with 16 GB of storage, can collectively pose as a cloud with a generous storage capacity of 128 GB.

Distributed storage for mobile ad hoc networks has been addressed by some research work (e.g.: [3], [4]), but most of them are not available as a concrete system for a network of smartphones, and [3] does not guarantee data persistence. In this context, our goal is to build a system that works out-of-the-box with current everyday mobile devices.

2. CHALLENGES

A distributed storage system for mobile ad hoc networks raises several issues that must be addressed. First, to ensure data persistence in a volatile environment, the system has to embed a replication mechanism. Several replication protocols for mobile ad hoc networks have been proposed, addressing one or more issues present in a mobile environment: energy consumption, network partition, or scalability.

.

Secondly, choices must be made regarding the network's structure. In an unstructured network, nodes do not know the contents' location. Therefore, a flooding mechanism is required to search for content. To ensure data persistence, a cyclic mechanism to check the number of replicas must be implemented. In a structured one (e.g.: Distributed Hash Table (DHT)), the nodes have a partial knowledge of the network, each node stores a part of the contents and knows where is the remain contents location. Although, a DHT has the additional cost of the network maintenance, we believe this cost is smaller than the required mechanisms to accomplish our goals in an unstructured network.

A third challenge is how to support device-to-device ad hoc communications. To solve this problem, we intend to use WiFi Direct communications (e.g: [5]).

Lastly, a device's resources must be carefully managed to remain available to other tasks, and reduce battery usage.

3. FEATURES & PROPERTIES

We propose a fully decentralized storage system for mobile devices that follows the classical key/value store model, where each key maps to one immutable value and serves as its identifier. The system supports the following operations:

`put`$\langle key, data\_item\rangle$ - stores the given content in the system, making it available to other nodes;

`get`$\langle key\rangle$ - looks in the data store for the item mapped by the given key, returning it if found;

`remove`$\langle key\rangle$ - removes the content mapped by the given key if the content was put by this node; and

`search`$\langle wildcard\rangle$ - returns all keys in the data store that comply to the given wildcard.

The system provides the two following properties:data persistence, and energy-awareness. Regarding data persistence, the system follows a *best effort* approach. Given the volatile nature of the environment we are considering, the effects of churn prevent the system from guaranteeing utterly persistence, and in our system a data item persists as long as all the nodes storing (replicating) it do not disconnect at the same time. Since we are considering mobile devices, an energy-friendly system is of the utmost importance.

4. ARCHITECTURE & IMPLEMENTATION

In order to achieve a fair design (in terms of resources consumed), we opted for a fully *decentralized and symmetric* architecture, i.e., all nodes have the same role and run the same services. As depicted in Fig. 1, each node has an architecture with the following components:

MOBIQUITOUS 2015, July 22-24, Coimbra, Portugal

DOI 10.4108/eai.22-7-2015.2260263

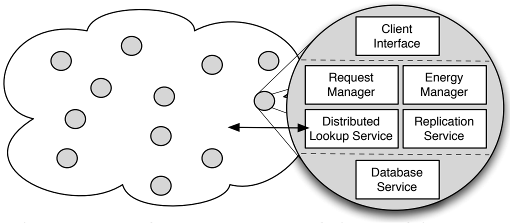

Figure 1: The system node's architecture

The Client Interface (CI). Provides the interface for user applications (`put`, `get`, `remove`, and `search`).

The Request Manager (RM). Coordinates the processing of the requests sent by the user applications: `get` operations are first directed to the local data layer (the *Database Service*), to check if they can be served locally and, if necessary, relayed to the *Distributed Lookup Service* for a network-wide search; `put` operations are handed by the *Replication Service* (RS) to send multiple replicas to remote nodes, and; `search` operations are relayed to both the local data layer and the DLS, for a complete coverage of the search space.

The Distributed Lookup Service (DLS). The purpose of the DLS is to track the whereabouts of the data items, and to perform all network-related management tasks. Our current prototype implementation resorts to TomP2P [1], a Java DHT that works on Android. In order to make TomP2P suitable for mobile devices and networks, we performed several optimizations, of which we highlight the reduction on the number of messages exchanged in: (i) the replication support mechanism, where before contents were proactively broadcast indefinitely to all the replicating nodes, and now content is propagated only when reacting to changes in the network; and (ii) the reduction of destinations to send messages for get and search operations. This approach is likely to increase the operation's latency but greatly reduces the number of exchanged messages and, consequently, the energy consumption in the mobile devices.

The Replication Service (RS). The RS is responsible for maintaining the number of replicas for each published data item according to a predefined replication strategy. The system grows from TomP2P and is completely decentralized, distributing the responsibility of guaranteeing the replication level of each individual data item by multiple nodes. For each data item, there is one node *responsible* for managing its replication in the distributed storage system, keeping track of the disconnection of replica-holding nodes. Whenever a replica-holding node disconnects from the system, a new one is selected and is sent the data item. The disconnection of a responsible node triggers the election of a new one by the replica-holders.

The system supports three different replication strategies. The *static replication* and *network-aware replication* are already included in TomP2P. The *replication by popularity* is new and works by counting the number of requests for each data item and increasing the number of replicas for the most popular data items. We expect this mechanism to reduce the latency of data obtainment operations and contribute to a more fair distribution of the energy consumption across the replica-holding nodes.

The Energy Manager (EM). When the battery level falls below a pre-established threshold the node is *partially* deactivated. The EM turns off the network-related maintenance tasks of the DLS, which results in making this node *invisible* to the remaining nodes in the system. In addition, the EM turns off the RS to stop the replication tasks.

The Database Service (DS). A key/value store to map the keys to shared and downloaded content's data.

5. CASE STUDY

As a case study, we designed an Android mobile application that mimics a shared photos gallery. Users may publish their own (local) photos, view a list of the photos available in the system, and download photos from remote nodes.

To assess the impact of our implementation in a smartphone's battery consumption, we conducted two preliminary experiments in a network comprised by two type of devices: Nexus 7 (2013) and Moto G 2nd gen. The first experiment targeted battery consumption when no operations, other than network maintenance, were being performed by the system. The results for an 8 hour run showed a consumption of 35% and 8% for, respectively, the device that created the network (i.e., the hotspot) and the devices that joined. An increase of 11% and 2%, respectively, when compared with the phone's execution without running the application.

The second experiment measured the amount of battery consumption on a node continuously requesting and downloading (and deleting) 2 MB images for an hour. In the end, the device downloaded and stored in disk 4962 images using 10% of battery. In turn, the node serving the requests presented a battery consumption of 14%.

6. CONCLUSIONS

This paper presented ongoing work on the development of a decentralized storage system for networks of hand-held devices. The preliminary experimental results show that our system may be used continuously during an event, such as a party, a concert or a sporting event, to share contents without draining all the battery.

Our current implementation needs an access point to communicate, such as a router or a mobile hotspot. Our next step is to enable ad hoc communication (e.g., Wi-Fi Direct) and, with that, extent the applicability of our system to scenarios where there is no wireless access point or even when the network is overloaded.

7. ACKNOWLEDGMENTS

This work was partially funded by FCT-MEC in the context of the research project CMUP-ERI/FIA/0048/2013, and under the research grant SFRH/BD/99486/2014.

8. REFERENCES

[1] T. Bocek. TomP2P. `http://www.tomp2p.net/`.

[2] DARPA Strategic Technology Office. Content-based mobile edge networking (CBMEN). `http://www.darpa.mil/program/content-based-mobile-edge-networking`.

[3] M. Lombera et al. Mobile ad-hoc search and retrieval in the iTrust over Wi-Fi Direct network. In *ICWMC*, 2013.

[4] J. Luo et al. PAN: Providing reliable storage in mobile ad hoc networks with probabilistic quorum systems. In *MobiHoc*, pages 1–12. ACM, 2003.

[5] A. Teófilo et al. Group-to-group bidirectional Wi-Fi Direct communication with two relay nodes. In *Mobiquitous*, 2015. to appear.